



HAL
open science

Semantic Frameworks to Enhance Situation Awareness for Defence and Security Applications

Valentina Dragos

► **To cite this version:**

Valentina Dragos. Semantic Frameworks to Enhance Situation Awareness for Defence and Security Applications. Information Retrieval [cs.IR]. Université de Paris, 2021. ⟨tel-03347626⟩

HAL Id: tel-03347626

<https://hal.science/tel-03347626v1>

Submitted on 17 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Semantic frameworks to enhance situation awareness for defence and security applications

Projet de Recherche

présenté et soutenu publiquement le 21 June 2021

pour l'obtention d'une

Habilitation à Diriger des Recherches

(Spécialité Informatique)

par

Valentina Dragos

Composition du jury

Président : Salima Benbernou, Professeur, Université de Paris

Rapporteurs : Iris Eshkol-Taravella, Professeur,
Université Paris Nanterre
Cecilia Zanni-Merk, Professeur,
Institut National des Sciences Appliquées (INSA) de Rouen
Aldo Napoli, Directeur de Recherche,
Centre de recherche sur les Risques et les Crises (CRC), Mines ParisTech

Examineurs : Joel Brynielsson, Directeur de Recherche,
Swedish Defence Research Agency (FOI), Stockholm
Simon Maskell, Professeur,
Université de Liverpool

Mis en page avec la classe thesul.

Sommaire

General introduction	5
1 Introduction and context of scientific contributions	7
1 Situation awareness : what is this about?	7
2 Technical bottlenecks of situation awareness	9
3 Knowledge and semantic technologies	10
4 The rationale for semantic-driven situation awareness	12
5 Structure of the document	13
2 Conceptual models to assess the quality of symbolic information	15
1 Motivations	15
2 The ONTO-CIF model and correlation of soft data	17
2.1 Credibility of information and ambiguities of textual data	17
2.2 Construction of the ONTO-CIF model	18
2.3 Correlation of HUMINT messages	21
3 Trust in reported information	23
3.1 Concepts and notions for trust assessment	23
3.2 A functional model of trust	25
3.3 Formalisation of the model and trust estimation	27
4 Related work	30
5 Lessons learnt	31
6 Collaborations and dissemination of results	32
3 Semantic integration in heterogeneous and dynamic environments	35
1 Motivations	35
2 Semantic interoperability for information indexing and retrieval	38
2.1 Knowledge acquisition for maritime surveillance	38

		1
2.2	Semantic integration of documents	38
2.3	Strategies for flexible information retrieval	42
3	Semantic mediation for on-the-fly integration of soft and sensor data	45
3.1	Dynamic fusion of sensor data and observations	45
3.2	Information extraction and semantic mediation	46
3.3	The INFOTEX prototype	51
4	Related work	52
5	Lessons learnt	54
6	Collaborations and dissemination of results	54
4	Content characterisation for cyberspace exploration	57
1	Motivations	57
2	Formal models for online content characterization	59
2.1	An ontology of appraisals for subjectivity modeling	59
2.2	Conceptual foundations of hate modeling in French	64
3	Illustration of ontology-driven social data analysis	69
3.1	Detection of trends in online data	69
4	Related work	75
5	Lessons learnt	76
6	Collaborations and dissemination of results	77
5	Conclusion and perspectives	81
	References	85

Acknowledgments

I would like to express my sincere gratitude to the committee of this research project: Prof. Salima Benbernou, Prof. Iris Eshkol-Taravella, Prof. Cecilia Zanni-Merk, Dr. Joel Brynielsson, Prof. Simon Maskell and Dr. Aldo Napoli. Thank you for accepting to review this report. It is an honor and privilege to defend my project in front of you.

Research is first and foremost a collective adventure and I am grateful to all the participants who contributed to my projects.

Many thanks to my fellow ETUR mates for the countless discussions we had and their continuous and contagious enthusiasm.

I am constantly learning from my students; let me express my gratitude to all the undergraduate students I supervised in the past, for their hard work and fresh ideas.

I would like to thank my lab mates, colleagues and research team for their wise counsel and sympathetic ear.

I also thank my family and friends for their unconditional support.

Credit goes to each and every one of you, who helped me in carrying out the best work pieces this report is based upon. All remaining errors and mistakes are mine.

General introduction

This document presents my research contributions that I gradually developed since joining ONERA, The French Aerospace Lab, in January 2010. Currently I am a member of the Department of Information Processing and Systems (DTIS), and from 2010 to 2017 I was a member of the Department of Information Processing and Modeling (DTIM).

My research was conducted in a department whose research activities address a large spectrum of information processing techniques, from logical reasoning, to operational research to multi agents systems. Applications are in three main fields: avionics, space, defence and security. With respect to the existing research directions, my research brought new insights on defence and security domains, by focusing on methods and techniques for *soft data* processing, which is to say symbolic data provided by humans (operators, Human Intelligence) or gleaned on the Internet (Open Intelligence, Social Intelligence). My research addressed the assessment of quality for human reported information, the characterization of online data, the development of semantic interoperability for heterogeneous information fusion and the analysis of uncertainty for information fusion.

This document reports results obtained from early 2010 to late 2020, as a research scientist at ONERA. Although my work was part of a new research line in my department, several projects such as ROSARIO (Research in Open Sources: information and opinions retrieval) and RIAD (Network on Artificial Intelligence and Decision Making) allowed me to collaborate with colleagues from my department, but the majority of my contributions are the results of external collaborations and projects.

I was the scientific coordinator of two research projects: FUTHANE (Fusion Technique de Haut Niveau et Evaluation), a project funded by the French Defence Procurement Agency to investigate the fusion of heterogeneous information, and FLYER (Intelligence Artificielle pour Analyser les Contenus Extrémistes sur Internet), a project funded by the French Research Agency thanks to its ASTRID (Accompagnement Spécifique des Travaux de Recherches et d’Innovation Défense) program. FLYER addresses the detection of extremist content online with artificial intelligence methods. Some of my scientific contributions were carried out within those project’s time frameworks and context.

Part of my work was also done in the context of several European and international projects : QUALIPSO, a joint research action France-Canada and three European projects: I2C (FP7, European Commission) ¹ , PRACTICIES (Partnership Against Violent Radicalisation in Cities) ² and AI4EU (Artificial Intelligence for European Union) ³

¹<https://cordis.europa.eu/project/id/242340>

²<https://practicies.org/>

³<https://www.ai4eu.eu/>

(H2020, European Commission).

Results reported in this document are also the fruit of several collaborations: at national level I had collaborations with partners from academia: MODYCO (Modèles, Dynamiques et Corpus UMR CNRS-Université Paris Nanterre), EFTS (Education, Formation, Travail, Savoir) Toulouse 2 University (Jean Jaures), LIDILEM (Laboratoire de Linguistique et Didactique des Langues Etrangères et Maternelles), University of Grenoble Alpes. I had joint projects and collaborations with industrial partners: Thales TRT, Airbus DS and Bertin IT.

At national level, some of my collaborations were facilitated by LEMON (anaLysE et dynaMique des messages et cONversations radicales sur Internet) ⁴, a coordination action supported by the French National Center for Scientific Research via MADICS (Masses de Données, Informations et Connaissances en Sciences), a research cluster promoting broad data-driven solutions. Together with Delphine Battistelli, full professor at Paris Nanterre University, we coordinate LEMON since 2017.

At international level, since 2011 I am a member of the ETUR (Evaluation of Techniques for Uncertainty Representation) working group, where I collaborate with researchers from George Mason University (Washington, USA), Thales TNO (Netherlands), University of Pretoria (South Africa) and IABG (Industrieanlagen-Betriebsgesellschaft mbH, Germany). In January 2018 I was appointed by the French Ministry of Defence to integrate the NATO Research Track *Exploration of Cyberspace for Intelligence* (IST 159) and this participation facilitated collaborations with FKIE (Fraunhofer Institute for Communication, Information Processing and Ergonomics, Germany), DRDC (Defence Research and Development Canada), NATO STRATCOM (Latvia) and the University of Helsinki (Finland).

I am also an expert for the French Research Agency and the European Commission for calls dedicated to security or information technology.

This document is a synthesis of my research project entitled *Semantic frameworks to enhance situation awareness for defence and security applications*. The first chapter introduces both the theoretical background and application context needed to understand the contributions presented in the rest of the document. The remaining chapters describe three semantic frameworks developed to improve situation awareness in various security of defence contexts: the evaluation of information items provided by human operators, the development of semantic-based approaches for information fusion in heterogeneous and dynamics environments, and the analysis and characterization of online data.

A semantic framework comprises both a knowledge model and the processing algorithms implemented upon, my research being at the crossroad of knowledge engineering and textual data mining.

Contributions are not presented chronologically, but rather regrouped according to their coherence, as several research lines were addressed at different moments in time and for different applications. Each contribution is discussed by indicating the research context, the maturity level of artefacts developed, main collaborations, lessons learnt and a selection of relevant publications.

⁴<http://www.madics.fr/actions/actions-en-cours/lemon/>

Chapter 1

Introduction and context of scientific contributions

Contents

1	Situation awareness : what is this about?	7
2	Technical bottlenecks of situation awareness	9
3	Knowledge and semantic technologies	10
4	The rationale for semantic-driven situation awareness	12
5	Structure of the document	13

This introductory chapter is threefold. First, we introduce situation awareness, a key concept for defence and security applications, and its main technical bottlenecks. Situation awareness is the application context of my work. Second, we review the rationale for using semantic frameworks to addresses technical bottlenecks of situation awareness. Finally, we briefly describe the remainder of this document.

1 Situation awareness : what is this about?

What is situation awareness?

Situation awareness is an application field which derives its research and technical findings from the proper management of critical situations in a variety of contexts. Complex crisis due to conflicts or natural disasters give rise to dynamic and unpredictable events: political transitions, insurgencies, terrorism, ethnic tensions or social upheaval. Responses may be build upon military, political, diplomatic, cultural or economic factors, but regardless of the nature of envisioned solutions, evidence is needed to inform decision makers.

Situation awareness can be defined simply as *knowing what is going on around us* or more technically as *the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning and the projection of their status in the near future* [52]. In practice, situations are simplified representations of complex constellation of actors, organizations, events and actions together with hypothesis about possible impacts.

Who needs situation awareness?

The need for situation awareness arises in a variety of environments: maritime security [91], human-machine interactions [61], autonomous driving [78], cyberspace [39], cybersecurity [44] or avionics [65]. Situation awareness is also crucial for defence and security applications.

How is situation awareness built?

Situation assessment combines information from multiples sources to reason about several entities over a range of time horizons.

From a technical standpoint, building situation awareness includes pulling together disparate data from numerous sources and correlating, fusing, merging and visualizing the data such that the end user can rapidly understand the situation and make decisions. Situation awareness requires systems and methods for information fusion, which further involves a process of combining data from multiple inputs, i. e. from one or more sources such as sensors and textual reports. The aim is to obtain information that is better, more useful and meaningful than that would be derived from each of the sources individually, that is, without fusing.

To avoid terminological confusion, in the context of this work the term entity refers to vehicles, persons, or convoys, organizations or communities in real or virtual environments.

Each entity is described as a vector of features, which, according to the sensor data used in the fusion process, provides the position and kinematics of the entity, its type and also relations to contextual information, such as geographical features (roads, airways) or to other entities in the situation. An entity is described as a set of states, representing the knowledge of this entity at any moment in time during the surveillance task. Entity state gathers the estimated features and additional information related to traceability and information assessment.

Let E_i be an entity, having a set of states ES_k , with $ES_k = (t_k, K_k, Tr_k, A_k)$ a time stamped vector of features, composed of the knowledge K_k including kinematics, nature and additional properties, the traceability Tr_k to observations used to produce K_k and the assessment A_k of K_k , represented as a probability, a likelihood or even as a simple confidence score. Entity states can be built upon sensor-based data and soft observation reports: this only depends on the ability of the algorithms to associate these observations with a given entity.

A situation of n entities is defined as the union of $E_{p,p \in (1..N)}$ the set if entities, the set of r sensor observations $O_{p,p \in (1..q)}^{sensor}$ and the set of q collected observations $O_{p,p \in (1..q)}^{soft}$, some of which could be false alarms, or inaccurate or misleading reports.

Who builds situation awareness?

From the very beginning, researchers on information processing have contributed to this topic, attracted by the challenging aspects of the problem. Although the need for automatic data processing for situation awareness has never ceased, nowadays it is ever increasing, triggered by the growing volumes of data available off and online, the avail-

ability of sophisticated communication means and the increasing complexity of defence and security domains.

The trend is not new and it is reflected in the establishment of proper conferences series (e.g., International Conference on Information Fusion, European Conference in Intelligence Security Informatics), dedicated topics within various European call for projects, international savant societies and organisations (ISIF- International Society for Information Fusion, European Defence Agency) and the establishment of specialized workshops, work groups or research tracks at international organizations. Examples are the ETUR (Evaluation of Techniques for Uncertainty Representation) working group supported by the ISIF (International society for information fusion) or the Inter-Group Workshop on *Big data challenges: Situation awareness and decision support* organized by NATO Science & Technology.

It is critical to observe that the entire panoply of methods, theories and knowledge that form the basis for situation awareness is provisional and complete knowledge of actions, crisis, events in real environments can never be achieved. Put simply, we cannot become and remain aware of everything going on around us. This is even truer in the case of defence and security applications, which often feature intelligent and dynamic adversaries who exploit unexpected opportunities. The complexity of environments and events makes the task of building reliable situation awareness particularly challenging.

2 Technical bottlenecks of situation awareness

Situation awareness, aiming at a comprehensive view of ongoing actions, inherently depends of data from a multitude of sources. Traditionally, situation awareness was supported by information fusion systems that handled sensor and soft data fusion as two distinct problem sets, based on the intuition that numerical sensor data seemed to be more suited for track-and-detect applications while human reports and open sources were rather requested to analyze asymmetric threats.

The following example shows that this sensor vs. soft data dichotomy is out of date, and makes the case for a unified approach. Let us consider a convoy of vehicles that illegally crosses the border between two countries. As with any illegal crossing, an alarm is triggered and the entity is tracked by a sensor fusion system. While the convoy approaches a city, one of the vehicles enters the urban area and losing its track is highly possible, as urban terrain has a very dense traffic and affects visibility and line-of-sight communications. Once track loss occurs, any eyewitness sighting or testimonies from bystanders on their day to day activities can be of interest and the track-and-detect problem requires not only sensor based fusion but rather a joint analysis of sensor and soft data. The task does not require qualified field analysts, but selecting people who are close to the incident and who may be opportunistic sources and provide meaningful input.

The problem is even more crucial when it comes to complex crisis in the society, whit social media adding to the problem and playing an active role in providing both relevant and irrelevant information in (almost) real time. As a consequence, the perception of this separate analysis of hard and soft information has changed: while early approaches concentrated on improving detection rates or accuracy of classification [158], current research

prone joint processing rather than specific procedures [83], tries to place the human back into the loop [121] and intertwines data integration and data analysis [58].

There are several technical bottlenecks arising thanks to this paradigm change, all related to the complex nature of symbolic information to be processed [115].

Assessing the quality of information items

Using imperfect and incomplete data requires procedures to assess the quality of information. Information evaluation becomes a critical aspect and a research topic on its own and the integration of symbolic data comes with specific challenges: items collected can be either partially or completely inconsistent and for operators it is then difficult to distinguish relevant data from large and noisy data sets [130].

Combining heterogeneous sources

Another critical aspect is the need to build the situation by accommodating heterogeneous inputs. Even if there is the capacity within a fusion system to properly handle sensor data, on one hand, and soft information on the other, determining the appropriate level of aggregation and setting appropriate integration frames are far from simplistic.

Making sense out of large streams of data

Another critical aspect is related to in-depth characterization of symbolic content and gaining insights from data collected over time. More specifically, procedures are needed to process social data, as the Internet is nowadays a prominent dissemination environment for sensitive content such as propaganda and extremist ideas.

Investigating the uncertainty

The last important critical aspects is the analysis of uncertainty. Defence and security applications take place in the ever-changing environment and require continuous monitoring and adjustment of the situation, either to keep it on track, or to recognize when a situation has changed sufficiently to require input adjustments.

This document outlines my contributions developed to address several technical bottlenecks of situation awareness by adopting semantic-driven solutions. Before explaining the structure of the document, the next paragraph introduces the rationale of using semantics in the context of situation awareness.

3 Knowledge and semantic technologies

Knowledge Technologies is a generic name for a variety of computer-based techniques, tools and methods that allows capturing, representing, retrieving and reusing domain-specific knowledge. Knowledge-based solutions start with the identification of knowledge artifacts that are important for decision makers and their activities, acquire the knowledge

and then represent this knowledge by using machine readable formalism in order to make it accessible to machine processing, easy to retrieve and reuse.

According to [138], the main fields of Knowledge Technologies are:

- Knowledge Based Systems (KBS) : those systems maintain a knowledge base storing a set of symbols of the computational model in form of statements about the domain, and they perform reasoning by manipulating those symbols.
- Natural Language Processing: tools and methods developed to extract meaning from unstructured or structured data in natural language;
- Data Mining: concerns the extraction of previously unknown knowledge from structured (databases) or unstructured data sets;
- Semantic Technologies: is a generic umbrella for a mix of methods from natural language processing, data mining and knowledge representation intended to derive meaning from information items; with other words, semantic technologies aim to make sense out of large or complex data collections;
- Case Based Reasoning: is a reasoning paradigm analysing a set of patterns or rules in previous situations (cases) so they can be applied to solve problems in new situations;
- Intelligent Agents: are computational systems that can sense and act in an autonomous manner in complex and dynamic environments;
- Document Management Systems: cover several techniques that support the life cycle of electronic documents and facilities the management of their contents and also their accessibility and retrieval.

My research work focuses on *semantic technologies* and their use to facilitate situation awareness in defence and security applications. Semantic technologies rely upon more complex models for knowledge representation: ontologies.

Ontologies are formal representations of knowledge and the most commonly adopted definition was proposed by Gruber [160] who defines an ontology as *a formal and explicit specification of a shared conceptualization* for a given domain. This definition highlights several characteristics of an ontology as a formal model of knowledge: having a formal representation, being explicit, capturing a shared conceptualization of domain-specific knowledge.

The first aspect is the formal representation and requires ontologies to be expressed in a knowledge representation language that provides formal semantics. Thanks to this principle, the knowledge represented by ontologies is machine readable and can be processed by automatic procedures.

The second aspect is the explicitness: ontologies describe domain knowledge in an explicit manner in order to make it accessible for machine processing. Notions, relations and associations that are not explicitly mentioned in the ontology are not part of the machine-interpretable conceptualization that the ontology captures, and are not available

for processing and inferences although humans might take them for granted in the light of common sense.

Conceptualization: refers to the ability of ontologies to specify the knowledge in a conceptual way, in terms of symbols representing concepts and their interactions. Thus, the ontology describes a conceptualization in general terms, instead of capturing specific states or situations, and tries to cover as many situations as possible.

Being shared: an ontology always reflects an agreement on domain conceptualization. This agreement is reached by different actors within a community and the resulting ontology is always limited to capturing the vision of this particular group of people in the community.

The last aspect of the definition refers to domain specificity and requires ontologies to be limited to describing knowledge about a particular domain of interest.

Ontologies describe domains of interest by using concepts, relations and instances. Concepts are ontological categories that are relevant in the domain of interest. Relations are connections of concepts and correspond to their interactions in the domain. Instances or individuals represent the named and concrete objects within the domain of interest. In summary, an ontology is a conceptual yet executable domain model that can be defined as:

$$O = \{C, R, H^C, H^R\} \quad (1)$$

where C is a set of concepts, R is a set of relations and H^C and H^R are hierarchies defining a partial order over the set of concepts and relations, respectively.

Simply put, every ontology is *a treaty, or a social agreement, among people with some common motive in sharing* [144]. Those models capture relationships and enable users to interlink multiple concepts in a variety of ways.

4 The rationale for semantic-driven situation awareness

Why using ontologies for situation awareness?

Semantic technologies are often used to gain an understanding of ill structured and wicked domains, because they allow for the collaborative structuring of field into agreed-upon concepts, attributes and relationships between the entities. Developing an ontology is akin to defining a set of knowledge and their structure for other programs to use. Semantic frameworks are understood as coherent structures of *ontologies and processing algorithms*.

Building semantic frameworks allows us to simplify the distribution, sharing and exploitation of information and knowledge, across multiple distributed actors, sources and decision makers. More specifically, ontology-based solutions are useful to solve several critical research aspects of situation awareness:

- Providing finer analysis of data quality. The quality of information can be analyzed in the light of a common understanding of the domain, achieved thanks to an ontology.

- Mastering the heterogeneity of contents. Heterogeneity within information fusion systems can be handled by using an ontology and its shared conceptualization as a mediator between information sources, users and domain experts.
- Performing in-depth analysis of large data sets. Large volumes of data can be characterized by using a set of concepts, instead of a set of words and inferences can be used to further refine the description of the content by aggregating those concepts.
- Coping with data scarcity. By their very nature, defence and security are data-scarce domains and involve applications where there is only a limited number of examples. For crisis management, for instance, accurately annotated data sets do not exist for many of the interesting activities one would wish a machine to detect: anomalies, suspicious behaviors, and other infrequently occurring activities. My approach using semantic frameworks is fundamentally grounded upon the idea of an alternative knowledge representation that removes the need to learn from important volumes of data.

5 Structure of the document

This thesis outlines my contributions developed to address several critical aspects of situation awareness. Results reported in this document were obtained from early 2010 to late 2020, as a research scientist at ONERA, the French Aerospace Lab, in Palaiseau. The documents is structured into there chapters.

Chapter 2 investigates several issues related to the assessment of quality of information items and describes two contributions: the first is the development of an ontology for information scoring and its exploitation to analyze messages provided by human operators in the field. The second contribution describes a model developed to estimate trust in reported information.

Chapter 3 presents two solutions developed to implement semantic interoperability in heterogeneous and dynamics environments. The first approach was developed to improve information indexing and retrieval for maritime situation awareness, while the second approach facilitates the fusion of soft and sensor data for dynamic situation assessment.

Chapter 4 addresses the analysis of data gleaned on social media, with emphasis on content characterisation. The chapter illustrates efforts towards providing conceptual foundations to detect features of subjectivity (emotions, opinions, beliefs, etc.) in online content and more specifically hateful messages.

The last chapter conclude this research project and outlines several perspectives for future work.

Chapter 2

Conceptual models to assess the quality of symbolic information

Contents

1	Motivations	15
2	The ONTO-CIF model and correlation of soft data	17
2.1	Credibility of information and ambiguities of textual data . . .	17
2.2	Construction of the ONTO-CIF model	18
2.3	Correlation of HUMINT messages	21
3	Trust in reported information	23
3.1	Concepts and notions for trust assessment	23
3.2	A functional model of trust	25
3.3	Formalisation of the model and trust estimation	27
4	Related work	30
5	Lessons learnt	31
6	Collaborations and dissemination of results	32

1 Motivations

The complex computational processes designed to build situation awareness can be represented as fusion systems, composed of numerous information sources, each with various parameters and properties. Taking into account the quality of information items is of paramount importance to assess the situation in a reliable manner. Assessing the quality of each intermediate and final items integrated into and generated throughout the fusion process is also necessary to allow the user to fully understand the situation under analysis.

Techniques for sensor data integration have been studied extensively by the scientific community [145], [54], [140]. In such projects, the aim is to merge data from multiple sensors in order to detect, identify or classify the entities in the situation, and from a

practical standpoint the fusion helps building a more accurate picture. The term *sensor* refers to a physical device, which measures the changes in physical quantity e.g. temperature, humidity, light intensity of the sample or surroundings. The quality of sensor data relates to the faults, such as outliers, bias, drifts, missing values, and uncertainty, which should be detected, quantified, removed or corrected [6].

In contrast to sensor data, what is generally called *soft information* by the scientific community are usually contents created by humans and include natural language messages, structured or semi structured messages, voice transcriptions, image exploitation products, open sources documents and textual content found on websites.

Soft items generally involve more degrees of freedom than sensor data and their quality is not necessarily related to errors, but to a number of specific factors: the accuracy and precision of information provided (is the message unclear or ambiguous?), the reliability of the source (did the source already provide credible items? do we know its intentions?) and the conditions of observability (did the human actually observe the content or he or she is just reporting the content or even recounting content reported by a third party?). Provenance of soft data is more coarse-grained, content is more heterogeneous and the structure of data cannot be reasoned about. Humans as sources may pass on hearsay, form opinions, analyze and interpret what they hear or see. If sensor data are *rows of numerical values affected by errors*, human messages are, according to Rein, a mix of *facts, conjectures, hearsays and lies* [22].

Based on such salient differences with sensor data, the problem of integrating soft data has been gaining interest since the early 2000s. Several research projects have started to consider the analysis of human sources [55], addressing the trust in reported information and the credibility of soft items [139] or investigating the nature of uncertainty in soft data [116]. In the same time this trend gave rise to the organization of special sessions [95] and panel discussions [127]. Several data sets for benchmark were also developed [133]. Those initiatives had an impact within the information fusion community and the discussions underlined the fact that assessing the quality of information items was not equally difficult [134] to be answered by the various teams and the sensor vs. soft data dichotomy was not interpreted the same way by the various sub communities [113].

More specifically, two aspects drawn attention. First, human messages were rather short and exhibited low semantics, and this made them unsuitable for processing with statistic-based methods. Techniques using additional knowledge were needed, in order to overcome this lack of intrinsic semantics. Second, as humans have the ability to distort the information, whether intended or not, taking into account not only the items, but also their propagation from one source to another before entering the processing chain was of paramount importance.

The goal of contributions discussed in this chapter is to address the conceptualisation of information quality for situation awareness. The first contribution describes the construction of a core ontology for intelligence analysis that is further used to investigate the correlation of items. While this contribution focuses on items, the second contribution takes into account also the source of information, and describes a model developed for trust assessment.

2 The ONTO-CIF model and correlation of soft data

2.1 Credibility of information and ambiguities of textual data

On this section we give an overview of the requirements that gave rise to the need for an ontology to support the assessment of information credibility. This next section explains how the model was constructed and validated.

For defence and security applications, the quality of information is assessed on two dimensions: source reliability and information credibility. NATO standards [150] instruct users to consider these components independently and to rate them on two separate scales. More specifically, the assessment of information credibility incorporates confirmation by other independent sources as a key determinant. When dealing with symbolic data, the confirmation of one piece of information by another is rarely complete, and for this reason the correlation is most commonly adopted as a measure expressing the degree of overlap of two information pieces.

Difficulties encountered when analyzing the credibility of soft items are due to the inherent ambiguities of textual contents. Auger and Roy [136] identify two types of ambiguities of textual data, based on the semiotic triangle proposed by Ogden and Richards [163]. This representation corresponds to a three-fold linguistic-semiotic system, see fig. 1, expressing connections between a symbol, also called *signifier*, carrying a meaning attached to a concrete object of the world.




Figure 1: Semiotic triangle, Ogden and Richards [163]

Linguistic ambiguities (blue line, fig. 1) appear as many associations exist between symbols and their attached meanings, and they are intrinsic to natural language.




Figure 2: Polysemy and homonymy

Therefore, polysemy, see fig. 2, is the capacity for a signifier to have different and non related meanings, while homonyms are words having different meanings while sharing the same spelling and pronunciation, see fig. 2. On the other hand, referential ambiguities (red line, fig. 1) are induced by socio-cultural aspects and are the consequence of differences arising as the items have an unique production context (including space, time, author, etc.,) and many possible interpretation contexts.

In additions to those aspects, solutions for symbolic data processing should take into account particularities of defence applications: divergences between doctrines, various security levels and cultural differences. Those constraints give rise to a need for explicit representation of the domain to be used as a basis to develop procedures for credibility assessment.

My contribution is based on the development of a formal ontology, called ONTO-CIF (Ontology for Credibility of Information and Fusion) in order to manage the ambiguities of textual data and to offer support when analysing the credibility of information items based on their correlation.

2.2 Construction of the ONTO-CIF model

The ONTO-CIF ontolgy was created thanks to a four-phases construction process shown in fig. 3.

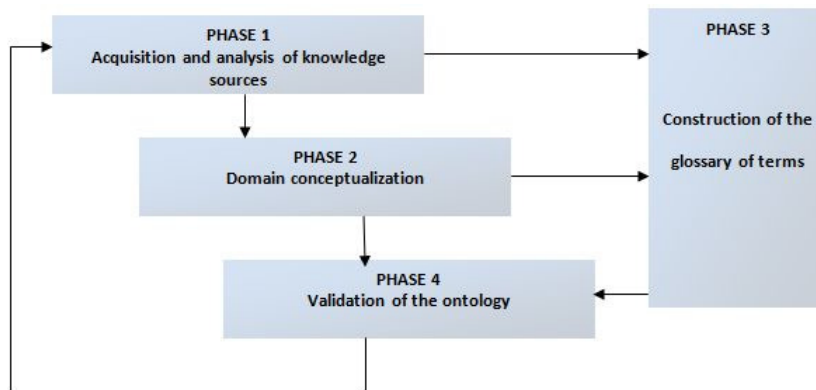


Figure 3: General approach to build ONTO-CIF

Phase 1: Acquisition and analysis of knowledge sources

This phase aims to create a collection of relevant knowledge sources to support ONTO-CIF development. For this work we used NATO standardization documents, in an effort to ensure the construction of a consensual representation shared by domain actors with respect to both conceptual and lexical points of view. In the light of those sources, a set of five notions corresponding to *Event*, *Place*, *Organization*, *Person* and *Equipment* (the pentagram, fig. 4) is the central element of the domain.




Figure 4: The pentagram

Those documents provide definitions of pentagram notions , i.e. an equipment is defined as *any item of materiel used to equip a person, organisation or place to fulfil its role*, and descriptions of entity associations.

The analysis of collected sources highlights several issues to take into account for the development of the ONTO-CIF ontology. Hence, there are no standardised taxonomies for military intelligence, although the AIntP-3 resource [151] was developed to facilitate exchanges of military data and NATO also maintains the APP6 glossary ⁵. As a consequence, those different attempts to conceptualize the domain introduce alternative and sometimes inconsistent definitions of the same concept. Moreover, the domain is modeled by considering different points of view: for instance, AIntP-3 provides functional descriptions of entities, while App6 is rather a collection of symbols used to mark military maps.

Phase 2: Construction of the glossary

During this phase we carried out the selection of domain specific terms. The glossary was created by including relevant terms naming concepts, instances, attributes or relations along with their synonyms. We restricted the construction of this glossary to the collection of documents and the glossary of terms is composed of 41 terms, see an excerpt in tab. 1.

Company	Social grouping	Terrorist attack
Flotilla	Chain of command	Entity
Civilian M	Event	Location

Table 1: Excerpt of the glossary of terms

Phase 3: Domain conceptualisation

The goal of this phase is to model domain concepts, to identify relations holding between them and to define formal axioms. The pentagram is the starting point of this phase,

⁵<http://www.nato.int/docu/stanag/aap006/aap6.htm>

therefore we create a basic ontological structure composed of: *Person*, *Organisation*, *Equipment*, *Locations* and *Event*. Then we gradually enrich this basic structure, by using the glossary of terms. For instance, *Geographical Area* and *Vehicle* are added as specific types of *Location*, and respectively *Equipment*. If the pentagram highlights concepts corresponding to the most relevant domain entities, those categories prove to be insufficient as soon as we are interested in describing the status of entities or their structures. To cope with those limitations, we introduce 3 abstract concepts, describing respectively: status of entities, according to various points of view (ex. *Military* vs. *Civilian*, *Natural* vs. *Artificial*) their structure (ex. *Hierarchy*) and their goals (ex. *Transportation*, *Communication*, etc.). During this stage we have identified 58 concepts, clustered in 6 main categories.

By taking into account temporal aspects, we introduce *Entity* as the most general concept regrouping time-independent artefacts and *Event* as the most general concept corresponding to some action that takes place at a given place and time, for instance *Terrorist Attack*.

Ontological relations are modeled by exploiting the set of associations identified during the previous phase. Those relations are: specializations, compositions and domain specific relations, see tab. 2.

Specialization	Organization, Social-Grouping
Composition	Organization, Person
Equivalence	Place, Installation

Table 2: Examples of relations

Axioms of ONTO-CIF define equivalent classes (*Installation = Place*), disjoint classes (*Military* vs. *Civilian*) and reversed roles (function-Of(Function, Installation) vs. has-Function (Installation, Function)).

Phase 4: Evaluation of the ONTO-CIF ontology

One of the challenges the ontology engineers have to deal with is their lack of knowledge as subject matter experts. To establish the quality of the ontology we adopted a procedure for ontology qualification and validation [143].

The qualification of the ontology is the process of establishing the relevance and adequacy of the abstraction of the real domain to the questions the model is expected to answer. This is a crucial aspect to be resolved during the conceptualization step, because the experiments designed to use the ontology for soft data analysis rely directly upon the capability of the ontology to accurately capture the domain. The qualification of the ontology was ensured by carefully selecting the sources describing the doctrine and technical procedures and also by iteratively co-creating the initial ontology : the selection of a set of concepts from the pentagram, the identification of their attributes and instances, the modeling of hierarchical and functional relations.

The validation of the ontology is the process of checking that the model covers relevant aspects of the real world on a credible manner. First, the ONTO-CIF model was validated

by bringing on board domain experts able to check the model from a basic inspection perspective. This involvement allowed for a fresh set of eyes to consider the ontology and allowed to refine namely the functional relations and the axioms. For instance, the bi-directional relations of the ontology were improved, and more specifically this allowed the modeling of inversed relations.

Independently, three criteria were selected from [68]. These criteria were used to analyze the completeness, the consistency and the redundancy of the model and are presented in tab. 3. Automated reasoning and expert analysis was used jointly to detect any logical inconsistencies in the ontology according to those criteria.

Criteria	Aspects verified	Analysis
Inconsistency	Circularity errors	No cycles detected by the reasoner
Incompleteness	Incomplete definitions	No missing concepts according to experts
Redundancy	Semantic redundancy	Axioms indicate equivalent classes

Table 3: Criteria to evaluate ONTO-CIF

Any inconsistencies were traced back and rectified in the light of expert insights and this step concluded the check for logical and domain consistency.

Abstraction level and practical use of the ONTO-CIF ontology

According to their abstraction level, ontologies are classified in three main categories: upper ontologies model high-level knowledge, based on philosophical reflections; core ontologies describe structural domain concepts and relations while domain ontologies provide detailed models of concepts as defined and handled by domain experts. With respect to the classification above, we have built ONTO-CIF as a core ontology, created to support intelligence analysis. It provides a generic description of the domain and the model had to be enriched for specific tasks. By choosing this abstraction level our intention is to manage the trade-off between the domain coverage and its detailed description, and to facilitate the exploitation of ONTO-CIF.

2.3 Correlation of HUMINT messages

The ONTO-CIF model previously described was used to estimate the correlation of HUMINT⁶ messages by adopting the principles of shallow semantic analysis (SSA). SSA was defined by Ferrandez and colleagues in [141], and consists in using ontologies to model domain knowledge, represent data as sets of ontological entities and then performing surface analysis on data having those semantic representations. The translation of free-form text into a semantic structure, normalized with respect to the ONTO-CIF model was carried out thanks to a preprocessing chain, including tokenizing, part-of-speech (POS) tagging, stemming and identification of named entities, as illustrated in fig. 5.

⁶HUMINT - HUMAN INTeLLIGENCE is a category of intelligence derived from information collected and provided by human sources

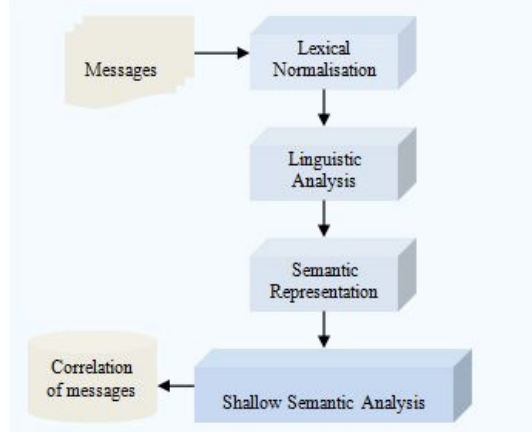


Figure 5: Processing pipe for shallow semantic analysis

Thanks to this phase, human messages m_i, m_j were further considered as collections of named entities E_i, E_j and ontological entities O_i, O_j . Based on this representation, the correlation of two messages was defined by taking into account two coefficients: the contextual and the semantic similarity, respectively. Therefore, the contextual similarity C is defined as the ratio between the set of named entities shared by both messages m_i, m_j , with respect to the overall set of named entities appearing in considered messages:

$$C = \frac{|E_i \cap E_j|}{|E_i \cup E_j|} \quad (1)$$

In a similar manner, the semantic similarity S is defined as the ratio between the set of ontological entities shared by both messages m_i, m_j , with respect to the overall set of ontological entities appearing in messages.

$$S = \frac{|O_i \cap O_j|}{|O_i \cup O_j|} \quad (2)$$

With those coefficients already defined, the correlation messages is then estimate as the sum of both contextual and semantic similarities.

$$Correlation(m_i, m_j) = S(m_i, m_j) + C(m_i, m_j) \quad (3)$$

This method was applied to estimate the correlation of HUMINT messages collected during a NATO tactical exercise. The corpus used for experimentation was composed of 81 messages reporting 6 different events, occurring during 3 days in different geographical regions. Messages were stored as XML files, having tags for several meta-data such as: geographical area of the reported event and the time of information delivery.

To assess the results of the correlation we have constituted a gold standard. We have conducted a study to collect manually assigned correlation ratings for selected pairs of messages. The study involved 5 researchers from ONERA and DRDC. The ratings were to be given consensually along a four step Likert scale [88] with the options: *strongly correlated*, *correlated*, *weakly correlated*, and *independent* plus an additional option *unsure*.

Unsure user ratings were not further considered in the evaluation. We asked users to perform their rating based on the context of messages and their meta-data: two messages are correlated if they refer to the same event. Additionally, we mapped the numerical values of the correlation algorithm to rating values as shown in tab. 4.

Correlation label	Numerical values
Strongly correlated	[0.85, 1]
Correlated	[0.75,0.85]
Weakly correlated	[0.65 0.75]
Independent	[0.00, 0.65]
Unsure	-

Table 4: Mapping of correlation values

On the basis of the expert annotations, we evaluate the correctness of the correlation algorithm, by comparing for each couple of messages the expert’s labels to the generated correlation values. This analysis allowed us to identify particular cases, when the estimation provided useful results, but also its limitations.

The approach performs well on identifying correlations of messages that share ontological and named entities and also identifies decorrelated messages. We also identified a few outliers, and namely the particular case of negations, when the algorithm estimates high correlation scores for contradictory messages while the experts rated *independent*. This result is counterintuitive although, from a correlation perspective, a message and its contrary are obviously highly correlated.

The main limitation of the solution is to perform only shallow analysis, and the results can clearly be improved by taking into accounts equivalence and disjunction axioms, in order to reinforce and, respectively, decrease the estimations of correlations. Another conclusion of this experimentation was that improvements of estimations can also be obtained by paying more attention to the set of distinct entities, in addition to considering only the set of shared entities.

3 Trust in reported information

3.1 Concepts and notions for trust assessment

The contributions above analyzed the intrinsic quality of information items, but from the perspective of an intelligence analyst, the quality of items should also include the impact of the source. This leads us to investigate the notion of trust in reported information.

To illustrate the problem analyzed in this chapter, let’s consider X , an analyst collecting evidence in order to decide whether or not an individual is involved in terrorist activities. In particular, he takes into account reports submitted by Y , a human source. Those reports usually consist on a mixed set of assertions (e.g., descriptions of events or states observed by Y) and opinions (i.e., judgments, assessments, or beliefs) expressed by

Y about assertion which give the analyst an insight into how strongly the source commits to the assertion, see fig. 6.




Figure 6: Assertions and opinions in human messages

In the statement contained in fig. 6, the source Y lets us know that she does not commit her full belief to the assertion that *John is a terrorist*, otherwise the reporter would have used phrasing such as *I am completely convinced* or *it is without doubt* or simply reported *John is a terrorist* as an unadorned statement. The information item is the sentence, which contains the assertion *John is a terrorist* and the uncertainty degree to be assigned because the analyst knows that Y is not completely certain about her own statements. The analyst must make a judgment about the veracity of John being a terrorist based upon factors such as previous experience with Y 's assessments in the past, or, perhaps, on the fact that other sources are relating the same information. Moreover, the goal of the analyst is to take this assertion into account, but also to encode his own belief about the quality of the source further in the analysis. All these different attitudes have to be evaluated by the analyst, who may have additional background information or prior evaluation of the source that have to be considered. At the end of this process, the analyst provides a level of *trust*.

Trust modeling is a research topic spanning diverse areas, from agent systems [87] to logical modeling to argumentation [100]. The concept of trust in these communities varies in how it is represented, computed and used. Although the definitions of trust vary from one domain to another, all those research areas consider trust as a user-centric notion that needs to be addressed in integrated human-machine environments as soon as those environments rely on information provided by external sources. Trust is a complex concept and several notions are relevant when considering trust analysis.

Trustworthiness of information sources is an attribute of sources that have the competences to report information, and who can be relied upon to share sincerely and clearly their beliefs on the uncertainty level of reported information. An item provided by such a source is then trusted by analysts.

Self-confidence [97] captures the explicit uncertainty assigned to reported assertions by the source. Statements may include the source's judgments when lacking complete certainty; these judgments are generally identified through the use of various lexical clues such as *possibly*, *probably*, *might be*, *it is unlikely*, *undoubtedly*, etc., all of which signal the source's confidence (or lack thereof) in the veracity of the information being conveyed. It should be noted that self-confidence, in our usage understood as the linguistic dimension of the certainty degree that the source assigns to reported items, is an aspect exhibited by the source, but it will be considered from the analyst's standpoint during trust analysis.

Reliability of sources indicates how strongly the analyst is willing to accept items from a given source at their face-value. Reliability is used as a general characterization that allows us to rate how much a source can be trusted with respect to their reputation, competence and supposed intentions.

Reputation of sources [148] captures a commonly accepted opinion about how the source performs when reporting information, and is generally understood as the degree to which prior historical reports have been consistent with fact. For human sources, reputation is considered by the analyst for each source based on previous interactions with the source and on the source's history of success and failure in delivering accurate information. Reputation relies, to a large extent, upon negative and positive experiences provided to the analyst by the source in the past.

Competence of sources [85] is related to a source's possession of the skills and knowledge in reporting on various topics: This aspect defines to what extent human sources understand the events they report on, whether the source has the ability to accurately describe those events, and how capable the source is of following the logic of processes producing the information.

Intentions correspond to specific attitudes toward the effect of one's actions or conduct. Reporting information can become *more a means to manipulate others than a means to inform them* [161] and thus can be carried out with the express purpose of inducing changes in another person's beliefs and understanding. Intentions are specific to human sources as only humans have the capacity to deliberately provide false or misleading information. Sensors may provide erroneous data because of device failures or environmental conditions, but never due to intention.

Based on those observations and notions, we have developed a conceptual model that estimates trust in reported information by taking into account the reliability of the source and the source's own characterization of reported items.

3.2 A functional model of trust

The model is composed of a source which provides an information item augmented with a degree of uncertainty captured by self-confidence to an analyst, as illustrated in fig. 7. Based upon his direct assessment of the reliability of the source, the analyst constructs his own estimation of trust in the item reported.

The model has two elements: the information source and the items reported from that source. The analyst is considered to be outside the model, although he or she can have multiple interactions with the elements of the model.

Definition of information source

An information source is an agent who provides an information item along with a characterization of its level of uncertainty. *Source* is a relative notion, depending on the perspective of analysis. In general, information is propagated within a chain relating real world information to some decision maker, and agents along the path can be both trained observers, whose job is to provide such reports, as well as witnesses or lay observers who may add items, in spite of not being primarily considered as information sources, but

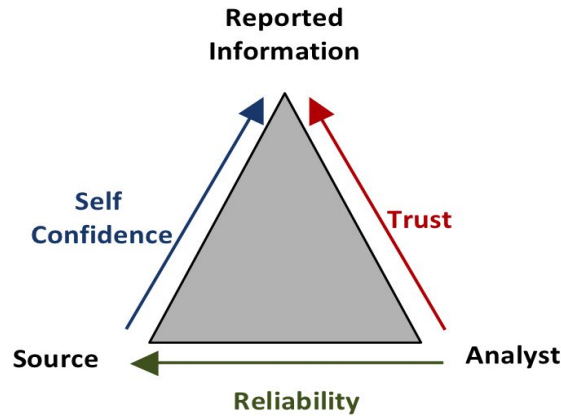


Figure 7: Model for trust analysis

rather as opportunistic ones. The notion of source is central in many information fusion applications and numerous research efforts aimed at modeling the properties of those applications. A general analysis of sources is undertaken by [129], who identify three main classes: S-Space, composed of physical sensors, H-Space for human observers and I-Space for open and archived data on the Internet.

Processing hard sensor information is widely covered [110] in the research community, and can be considered quite mature, while the integration of human sources brings many new challenges. Our model addresses human sources, and reported items can refer to actions, events, persons or locations of interest.

Information reported by humans is unstructured, vague, ambiguous and subjective, and thus is often contrasted with information coming from physical sensors, described as structured, quantitative and objective. While humans can deliberately change the information or even lie, sensors are also prone to errors and therefore hard information items are not always accurate. For human agents, the source is part of the real world, (a community, a scene, an event) and can be either directly involved in the events reported, or just serving as a witness.

Definition of reported information

Reported information is a couple $(I, \chi(I))$, where I is an item of information and $\chi(I)$ the confidence level as assigned by the source. Items are information pieces that can be extracted from natural language sentences, although the extraction and separation from subjective content are out of the scope for the model developed. Each item I has assertive i_a and subjective i_s components conveying factual and subjective contents respectively.

The model introduces two functions to estimate the reliability of sources and the self-confidence of information, as described hereafter.

Definition of a reliability function

A reliability function is a mapping which assigns a real value to an information source. This real value is a quantitative characterization of the source, inferred with respect to the source's previous failures, its reputation and the relevance of its skills for specific domains. For this model, the reliability of human sources combines three features: competence, reputation and intention. Competence captures the intuition that the quality of information reported by a source depends on the level of training and expertise, which may be designated as satisfactory or not, depending upon the task. Reputation is the overall quality of a source, estimated by examination of the history of its previous failures. Intentions refer to attitudes or purposes, often defined with respect to a hidden purpose or plan to achieve.

Reliability is a complex concept and, from a practical standpoint, it is difficult to have complete information about the global reliability of a source. Thus, this model describes reliability along the three attributes (competence of a source, its reputation and its intentions) described above. In practical applications, this solution allows for compensation for insufficient information on one or several aspects of reliability and to conduct, if necessary, the analysis of reliability based on just one attribute.

Definition of a self-confidence function

The self-confidence function is a mapping linking a real value and an information item. The real value is a measure of the information credibility as evaluated by the sensor itself and is of particular interest for human sources, as often such sources provide their own assessments of the information conveyed. Identifying features of self-confidence requires methods related to a research task of natural language processing: the identification of assertions and opinions in texts. In this field, the commonly adopted separation of those notions considers assertions as statements that can be proven true or false, while opinions are hypotheses, assumptions and theories based on someone's thoughts and feelings and cannot be proven.

The model keeps the distinction between sources and items in order to better dissociate the impact of both beliefs of sources and opinions of analysts on the source on the information provided. Using the model to assess trust values requires the estimation of functions and also a formal representation of the model. The rationale behind this is the observation that even reliable sources can sometimes provide inaccurate or imprecise information from one report to another, which is even more plausible in the case of human sources.

3.3 Formalisation of the model and trust estimation

The model developed combines various attributes of the source (discussed previously under "reliability") with "self-confidence" in order to capture trust of information as conveyed by the human. The model is source-centric, see fig. 8, and considers the source's ability to correct, alter or qualify the information report.

Although the rules for ranking, prioritizing and combining the attributes introduced




Figure 8: Trust estimation from source to analyst

by the model can be drafted empirically, the estimation of a trust value requires a formal representation of the model.

Evaluation of reliability

Assessing reliability is of real interest when opportunistic sources are considered because the analyst has neither an indication of how the source might behave nor the ability to monitor or control either the human providing the information or the environment in which the source operates. Various methods can be developed to estimate competence, reputation and intentions of the source. For example, competence is closely related to the level of training of an observer or can be defined by domain knowledge. Values can be expressed either in a linguistic form (*bad, good, fair, unknown*) or by a number. Reputation is an attribute which can be constructed not just by examining previous failures of the source but also by considering its level of conflict with other sources; this too can be expressed by numeric or symbolic values.

While reputation and competence can be, at least in some cases, estimated from prior knowledge, characterizing the intentions of a source is subject to human perception and analysis. Judgment of human experts is needed not just because there usually is no *a priori* characterization of the source with respect to its intentions but also because it is important to assess those aspects from the subjective point of view of an expert in the form of binary values only.

From a practical standpoint, it is suitable to provide an expert with a description of source competence, reputation and intentions as assessed independently. This way, experts can have the opportunity to develop different strategies of using reliability: they can decide to assign different importance to those attributes under different contexts or can use their own hierarchy of attributes. For instance, an expert may consider as irrelevant the information provided by a source whose competences is lower than a specific threshold or if he suspects the source of having malicious intentions.

A possible solution for estimating a unified value for trust is to consider reliability and self-confidence within the framework of an uncertainty theory and to rely on the set of combination rules the theory defines for example, those developed in probability theory, in possibility theory, or in belief functions theory. All these theories provide various operators to combine reliability and self-confidence in order to estimate trust.

Evaluation of self-confidence

Estimation of self-confidence aims at assigning a numerical value which captures how strongly the author stands behind assertions in the statement, on the basis of lexical clues he has included in the utterance. More generally, markers of an author's commitment are in the form of hedges, modal verbs and forms of passive/active language. A hedge is a mitigating word that modifies the commitment to the truth of propositions, i.e., certainly, possibly. Its impact can be magnified by a booster (*highly likely*) or weakened by a downtoner (*rather certain*).

Modal verbs indicate if something is plausible, possible, or certain (*John could be a terrorist, you might be wrong*). Moreover, in some domains sentences making use of the passive voice are considered as an indicator of uncertainty, in the sense that author seeks to distance himself from the assertions in the items reported through use of passive voice. Quantifying self-confidence is a topic of particular interest for intelligence analysis, and it was early addressed by Kent in 1962, [162] who created a standardized list of words of estimative probability which were widely used by intelligence analysts. This list has continued to be a common basis to be used by analysts to produce uncertainty assessments. Kesselman describes in [137] a study conducted to analyze the way the list was used by analysts over the past, and identifies new trends to convey estimations and proposes a new list having the verb as a central element. Given the variety of linguistic markers for uncertainty, the estimation of a numerical value based on every possible combination seems unrealistic, as the same sentence often contains not just one but multiple expressions of uncertainty. Additionally, assigning numerical values to lexical expressions is not an intuitive task, and Rein shows in [76] that there are no universal values to be associated in a unique manner to hedges or other uncertainty markers. As the author argues further, it is, however, possible to order those expressions and use this relative ordering as a more robust way to compare combinations of uncertainty expressions, and thus highlight different levels of uncertainty in natural language statements.

Formalization of the model and experiments

The aim of formalization is to provide a formal representation of the model, combining the capability to exploit the structure and relationship of elements of the model with the ability to express degrees of uncertainty about those elements.

Of particular interest is the observation that the developed model introduces a cognitive view of trust as a complex structure of beliefs that are influenced by the individual's opinions about certain features and elements, including their own stances. Such a structure of beliefs determines various degrees of trust, which are based on personal choices made by analyst, on the one hand, and the source, on the other hand. Therefore, the formalization requires a formalism that is more general than probability measures or fuzzy category representation, which are more suitable for applications considering trust in the context of interactions between agents. Moreover, the limitations of using subjective probabilities to formalize trust from this cognitive standpoint are clearly stated in [153].

As a result, the model was represented with belief functions, a formalism that is consistent with the cognitive perspective of trust adopted by the model. This belief-

based representation provides the most direct correspondence with elements of the model and their underlying uncertainty, while being able to quantify subjective judgments.

The formalization of the model with belief functions was conducted within the context of the ETUR working group. This formal representation was then used to perform trust assessment for several practical scenarios, involving reliable, incompetent or misleading sources, respectively. We report in [33] technical details of the formalization and experimentation.

4 Related work

Ontologies attracted much attention for defence and security applications, triggered by the rise of new processing techniques allowing for the domain knowledge to be incorporated within different systems [20].

For situation awareness, one of the first models developed to support situation awareness is the SAW (Situation Awareness) ontology [149], that models a situation as being composed of entities, interacting in order to achieve one or several goals. CONON (Context ONtology) [147] was another ontology having the context as central element, where context is an aggregation of four features: local coordinates, persons, objects and activities.

More recently, several approaches have considered the use of ontologies to investigate the problem of sense making from heterogeneous data [15], while other studies focused on providing solutions to manage the huge amount of information gleaned nowadays for intelligence purposes [104]. Techniques and methods are strongly correlated to concrete applications ranging from digital forensics [8] to crisis management [77] to investigation of criminal activities [25] and to governance, leadership and decision support [38].

From a practical perspective, ontologies are used as common ground to facilitate interactions between entities, such as UAV's and their ambient [18] or autonomous vehicles [81] or to build the semantic backbone for the Internet of Things (IoT) [5]. Ontologies are also used to investigate the quality of information, addressing features of veracity [9] or the uncertainty of relations [56], and knowledge model are core to generating social science inferences [30] for human-machine environments [23].

Internally, the ONTO-CIF model was also used for a different task of the QUALIPSO project intended to augment sensor data with semantics in order to facilitate information exchange for heterogeneous sensor networks [124].

Regarding trust assessment, the literature shows that this topic is currently part of a more general research line, tackling the quality of information used for intelligence purposes [28] and the design of standardized models for this task [16]. Indirectly, the model of trust propagation is also correlated to aspects of privacy [27], since attributes related to private matters (beliefs, opinions, convictions) are (sometimes) embedded in trust estimation.

Within the context of trust estimation for reported information, applications observe a shift from analyzing natural language documents [22] to social data analysis [7]. Contributions focuses namely on providing metrics [96] to estimate trust when information is distorted in space and time [29], or integrated within specific contexts [24].

For my research activities, the trust model was formalized with belief functions [10] and to what extent this formalization was in line with the elements of this model was further analyzed with the members of the ETUR working group on uncertainty analysis [33].

5 Lessons learnt

Assessing the quality of human messages was at the center of contributions discussed in this chapter. I showed that human messages have not only intrinsic but also extrinsic uncertainty and I made the case for using not only information extracted from messages (named entities, relations) but also domain knowledge when analyzing this type of data.

Initial approaches developed to integrate HUMINT messages supposed the transformation of natural language sentences into logical propositions. Converting sentences to their logical forms, going from words and their meaning to logical variables and their values, induced an important information loss. Those methods were purely logical, brittle and not robust enough to perform well on unrestricted naturally occurring text. Our purpose in modeling the ONTO-CIF ontology was to highlight the determinants of information quality for intelligence analysis and to depict the interconnectedness of information characteristics. This representation helped the analysis of HUMINT messages. Currently the analysis of messages in their textual form is still an active research direction.

The model for trust assessment was inspired by several studies undertaken by intelligence analysts that support the introduction of a single measure of perceived information quality (i.e., accuracy/truthfulness) incorporating all available attributes, including source reliability and information credibility [118].

From a practical perspective, the complete automation of symbolic information evaluation may be undesirable, given the requirement for analysts to easily understand and modify their inputs as new information becomes available. For example, the current emphasis placed on confirmation/consistency may reinforce order effects, given that new information must conform to prior information to be deemed credible. Also, when receiving second-hand information from a HUMINT source, one might consider the reliability of the primary source, the reliability of the secondary/relaying source(s), the reliability of the collector, as well as the reliability of any agent transmitting the information [111]. While information evaluation is largely accepted as a key function within the intelligence process, there are still currently open questions and research directions to be addressed by domain experts and researchers alike.

Although this research line started with HUMINT messages, the problem of how human communications are uncertain, both within the content and about the content, is still a central pillar of my activities, and the exploitation of open sources addressed in chapter 4 gave me the opportunity to investigate other aspects of this problem.

Today's challenges in human provided information for defence and security applications are still numerous. One of the most important challenges for end-users and practitioners alike is to have access to systems equipped with fusion modules able to deal with large-scale volumes of data. The need is in line with the rise of BigData paradigm and the prevalence of Internet as a major information source for intelligence applications [79].

It should be noted that, given the sensitivity of those domains, systems are not freely available to the scientific community and restrictions apply even to research prototypes.

6 Collaborations and dissemination of results

QUALIPSO - Joint research action France - Canada

Assessing the quality of information was the main research direction of QUALIPSO, a joint research action between the France and Canada funded by the French MoD.

The ONTO-CIF ontology was built from scratch and then used in the context of QUALIPSO project for empirical validation, by adopting the task of soft data correlation. I created the ONTO-CIF model, and also designed and implemented the algorithm for data correlation. I participated in experimental settings of the project to demonstrate the ability of this approach to deal with concrete scenarios. I also investigated different aspects to be considered when assessing HUMINT messages, such as contradictions [50], [117]. Experimentation was conducted together with **Anne-Claire Bourry-Brisset**, **Anne-Laure Jouselme** and **Patrick Maupin** [124] from DRDC, our Canadian partner for QUALIPSO. The ONTO-CIF model was of interest for practitioners in the field of the intelligence analysis and the model was presented in the frame of seminars organized by the French MoD [106] and French intelligence services.

Evaluation of techniques for uncertainty representations (ETUR) WG

Trust in reported information was a topic of the ETUR working group, to which I contribute since 2012. Several discussions with **Kellyn Rein** and **Joachim Biermann** [72] from FKIE, Germany, helped me clarify the need to analyze jointly the information and its source, as a first step of tackling the distortion of information throughout the reporting chain.

The model for trust estimation was developed with **Kellyn Rein** and later my colleague **Jean Dezert**, contributed to the formalization of the model with belief functions. More specifically, information credibility and features of self confidence were discussed with **Paulo Costa** and **Kathryn Laskey** from George Mason University, Washington and **Erick Blasch** from Air Force Research Laboratory, USA [103]. The ETUR working groups also offered me the opportunity to discuss the quality of knowledge models, with **Alta de Waal** and **Pieter de Villiers** from University of Pretoria, in South Africa and **Jürgen Ziegler** from IABG, Germany [11].

Supervision of students

Within this context I also supervised **Nezha El Fakraoui**, a student from the Intelligent Systems and Robotics Master at Paris 6. During her 6 months internship, Nezha correlated the assessment of soft data to uncertainty criteria developed by the ETUR working group. Nezha analysed which URREF criteria are relevant for soft data analysis and also identified the aspects uncovered by the URREF uncertainty framework. Her findings were also useful to consolidate this framework.

Dissemination of results

Dissemination of activities presented in this chapter include:

- Two journal articles published in the *International Journal of Knowledge-Based and Intelligent Engineering Systems* [107] and the *Journal of Advances in Information Fusion* [31];
- Organisation of the COTA workshop (Cotation des informations; Théorie et Applications) in conjunction with IC 2010 (Les journées francophones d'Ingénierie des Connaissances), a major national conference on knowledge engineering;
- Participation to the panel "On credible items, reliable sources and useless evidence", panel of the 16th International Conference on Information Fusion, FUSION 2013, Istanbul, Turkey.

Main publications of the chapter

- [10] Valentina Dragos, Jean Dezert, and Kellyn Rein. "Assessment of Trust in Opportunistic Reporting using Belief Functions". In: *22nd International Conference on Information Fusion, Ottawa, Canada*. IEEE. 2019, pp. 208–216.
- [33] Valentina Dragos, Jean Dezert, and Kellyn Rein. "Tracking uncertainty propagation from model to formalization: Illustration on trust assessment". In: *Journal of Advances in Information Fusion* 13 (2018), pp. 198–216.
- [103] Erik Blasch et al. "URREF reliability versus credibility in information fusion (STANAG 2511)". In: *Information Fusion (FUSION), 2013 16th International Conference on*. IEEE. 2013, pp. 1600–1607.
- [107] Valentina Dragos. "Developing a core ontology to improve military intelligence analysis". In: *International Journal of Knowledge-based and Intelligent Engineering Systems* 17.1 (2013), pp. 29–36.
- [117] Valentina Dragos. "Shallow semantic analysis to estimate HUMINT correlation". In: *2012 15th International Conference on Information Fusion*. IEEE. 2012, pp. 2293–2300.

Chapter 3

Semantic integration in heterogeneous and dynamic environments

Contents

1	Motivations	35
2	Semantic interoperability for information indexing and retrieval	38
2.1	Knowledge acquisition for maritime surveillance	38
2.2	Semantic integration of documents	38
2.3	Strategies for flexible information retrieval	42
3	Semantic mediation for on-the-fly integration of soft and sensor data	45
3.1	Dynamic fusion of sensor data and observations	45
3.2	Information extraction and semantic mediation	46
3.3	The INFOTEX prototype	51
4	Related work	52
5	Lessons learnt	54
6	Collaborations and dissemination of results	54

1 Motivations

In the previous chapter the problem of analyzing the quality of information items was addressed. In practical applications, situations are built by integrating not only several items, but also several types of information, provided by different sensors or human sources or a mix of both. This chapter is dedicated to two semantic-based solutions developed to assess situations in heterogeneous and dynamic environments. Those environments are characterized by the use of distributed, autonomous, diverse and dynamic information

sources, and the need to integrate numerical and soft data in dynamic and asynchronous fashions.

As an example of heterogeneous situation, consider a maritime scenario when several analysts are in charge of the surveillance of a maritime zone. The aim is to detect abnormal behaviors of vessels, and then to infer a category of those abnormal behaviors, among which illegal trafficking and smuggling are rather dangerous and a technical dysfunction is not harmful. The user can analyze the cinematic features of the vessel, can interrogate data bases, and complete the situation thanks to reports or items gleaned from the Internet. For example, the following steps can be followed: if the name or type of the vessel are identified, they are searched in domain repositories and the situation is completed with additional information. If operators can approach the vessel, in flight for example, their observations can be integrated. Inputs from those sources are then mixed up to build the situation, and in the light of this augmented description a decision as to yes or no the behavior is dangerous can be made. Results of this process, along with annotations, can eventually be uploaded in a knowledge base, to create a basis of cases already investigated for further analysis. While the example above illustrates the problem of heterogeneous fusion, dynamics refers to the ability to integrate information items on-the-fly. Managing dynamic situations requires the ability to make meaningful correlations of entities in the situation and incoming items.

Semantic integration has been identified as a key challenge, critical to implement successful assessment of heterogeneous and dynamic situations. In this context, the notion of semantic integration has a simple, appealing interpretation: adopting semantic standards such as knowledge models or terminologies, that support information sharing and fusion among technical bricks. In other words, items (messages, images, etc.) need to be standardized in order to be interoperable and used by several actor (humans and machines) in contexts different from the original one. Semantic integration permits the independence with respect to the collecting step (source, type, etc.) or the data processing context (information retrieval from repositories, fusion of items, data mining).

The need for semantic integration is not specific to defence and security domains, and for example in the medical domain Dolin and Alschuler stated already in 2011 that *challenges of semantic-less communication with today's model and terminology include ambiguities, lack of complete expressivity, redundant representations that cannot be computationally converted into a common canonical form, implicit semantics, and a less-than-perfect understanding of context* [122].

For defence and security domains, the emergence of operational solutions for semantic integration is hampered by the inability of applications to conform to semantic standards. First, these applications provide interfaces to security professionals in order to collect data in a way adapted to their use and suitable for their daily practice but usually not conform to standards and not intended for sharing. Then, in order to collect information in an evolutionary manner, security applications are often based on information models that are locally implemented. Finally, the principles of structuring and coding information and the level of granularity of information can also vary depending on the security professionals or the activity mode (surveillance, crisis management, intelligence analysis, etc.).

Contributions described in this chapter consider semantic integration along two directions: semantic mediation and semantic interoperability, see fig. 1.




Figure 1: Semantic interoperability and mediation

Semantic mediation allows the integration of disparate data items. More particularly, ontologies are used in order to combine items provided by different sources or results provided by analytical components. Items are not linked or connected, but their joint analysis and integration into the situational pictures is driven by an ontology.

Semantic interoperability is understood as a way to provide linkage, integration and meaningful use of data between users, experts and systems. More particularly, ontologies are used as a common and shared ground to enrich data with context and meaning, and also to facilitate the retrieval of information.

Our contributions are the following: the first contribution introduces semantic interoperability for heterogeneous maritime situation assessment. This solution uses an ontology in order to achieve enhanced understanding and better evidence-based interpretation of situations for vessels and maritime traffic within the context of maritime surveillance. The second contribution relies on semantic mediation to combine human observations and sensor data in the context of dynamic situation assessment. This contribution uses ontologies in order to create a bridge between two complementary representations of a situation that changes over time.

For both contributions, the use of ontologies has a dual objective: (i) reduce the structural complexity of resources and to make resources, in our case documents, looking simpler, easy to “grasp” by users, and (ii) make the complexity of dynamic integration of resources lower, because they are performed on simpler structures having eventually shared concepts.

2 Semantic interoperability for information indexing and retrieval

This contribution was developed in the context of the I2C European project and its goal was to illustrate how semantic integration can improve the reuse of previous analysis carried out by domain experts.

2.1 Knowledge acquisition for maritime surveillance

Surveillance in the maritime domain is undertaken by operators and analysts around the world to identify and assess the risk of threats and illegal activities. A typical surveillance scenario involves operators analyzing facts from a picture automatically build by fusing video, radars, AIS (Automatic Identification System) and intelligence information. Beyond analyzing the current maritime situation, this human-centered task also leads to the creation of *interpretation reports*, providing details on the suspicious behavior of vessels and conclusions of investigation. Created by domain experts, interpretation reports are the main source of knowledge for future investigations. The reports describe several types of suspicious behaviors but also the set of conditions to be satisfied in order to notify them as worthy of interests. In practice, experts create huge collections, containing hundreds of reports. Searching relevant reports in those collections becomes then a daunting task for humans: searching for information is tedious and error-prone for users because one must separately probe each source, verify semantic relevance to the new incident, and then manually integrate the results.

The problem addressed in this section is how to improve the access to the collection of interpretation reports. The solution consists in an information navigation model that was developed to explore the collection of reports and to guide the user towards documents that are relevant to its information need. A domain ontology supports the overall solution : first, the collection is indexed based on the ontology and then user queries are build by using this ontology. Thanks to those ontology-driven procedures, semantic interoperability is at the core of the navigation model.

2.2 Semantic integration of documents

Semantic integration is carried out by building an index to structure the collection of interpretation reports. Within the context of this work, a main difficulty was to build an index whose structure offers a multi-faceted description of the collection of documents and to create flexible information retrieval strategies based upon.

Construction of interpretation reports

Reports are created by a semi-automatic procedure, as shown in fig. 2.

First, a tool developed by partners of the I2C project performs real-time analysis of maritime situations and builds a draft of the interpretation report. The maritime situation is built from multisource information fusion, such as shoreline sensor data or deployed platforms. A learning engine and a rule engine access and analyze the traffic

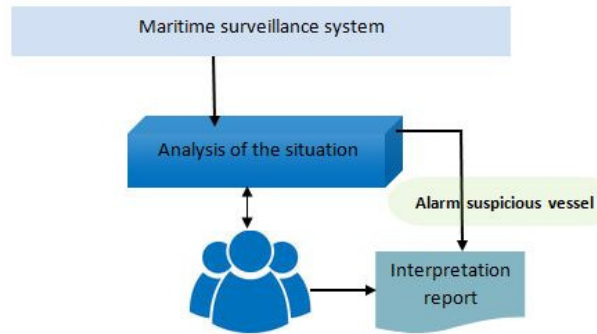


Figure 2: Semi-automatic construction of interpretation reports

picture in order to detect clues of abnormal behaviors such as deviations from standard routes, unexpected low velocity or AIS activity. If such activity is detected, an alert is triggered in order to draw the attention of operator towards vessels that are worthy of interest. An investigation is undertaken only for vessel of interest, in order to identify the nature of threat associated. A draft of the interpretation report is then created to gather data about the suspicious behavior and the conclusions of the investigation. Second, domain experts improve the draft by adding natural language paragraphs to explain the event under analysis. Reports are created semi-automatically as the system is able to provide a set of data that help operators to perform further investigations.

According to Waltz’s terminology introduced in [152] interpretation reports are composed of data and knowledge units. Data units are observables whose values are gathered through observations, and knowledge units are indicators and hypothesis, acquired thanks to domain knowledge and expert’s interventions.

Observables are a set of variables describing the current situation such as: details about vessel (name, type, activity, information about the crew, cargo and owner), the spatial-temporal context (geographical coordinates, time of alert) or meteorological conditions (weather and navigation).

Indicators are conditions of suspicion expressing signatures of threatening behavior. For instance, vessel older than 15 years or vessel stopping during night time can be relevant for abnormal or threatening comportment. Vessels can be also flagged if they start the itinerary in a zone considered as critical.

Hypotheses are models of specific abnormal behaviors automatically associated by the system to vessels under analysis. Illegal immigration, pollution or traffic are of particular interest for maritime safety experts.

While threat analysis is supported by combining observables, indicators and hypotheses, the last ones are critical information pieces, as potentially they provide explanations as to why the comportment may be of interest.

A hybrid index for information indexing and retrieval

In order to organize documents in the collection and to facilitate their retrieval, we developed an index able to connect observables, indicators and hypothesis, as shown in fig.

3.




Figure 3: Data and knowledge layers of the index

The index has two levels: the data layer, intended to handle numerical values of observable, and the knowledge layer, intended to handle domain concepts of indicators and hypotheses. The data layer is a basic tree-based index while the knowledge layer is an ontology of the maritime domain.

To capture the maritime domain and assist in the situational awareness, we created an ontology that describes the attributes and behaviors of maritime vessels transporting legally and illicitly goods and persons across the sea. This ontology was built from scratch, within the framework of the I2C project, and the construction involved a team of maritime safety experts involved in the project. Fig. 4 illustrates the upper level of this ontology.



Figure 4: Ontology of the maritime domain

Besides modeling domain knowledge, the ontology takes into account the structure of documents, as *Observable*, *Hypothesis* and *Indicators* are central concepts of the ontology.

The data and ontological layers of the index are connected through *instance-of* or instantiation links. This hybrid index ensures a twofold description of the collection and its structure allows us to carry out traditional information retrieval strategies, restricted on data level, on one hand, and advanced information retrieval strategies, corresponding to enhanced user information needs, on the other.

From interpretation reports to abstract representations

Each interpretation report is composed of several homogeneous semantic units. From a structural point of view, those units correspond to XML tags and describe: observables and their values (attribute-values sets), conditions of vessel suspicion (semi-structured data containing indicators), hypotheses raised by the system along with corresponding arguments written by experts. By taking into account both the content and the structure of reports, the model adopted to represent each document is a triplet:

$$D = (V_d, C_d, W_d) \quad (1)$$

where : $V_d = (v_1, v_2, \dots, v_n)$ is the vector of numerical values of observables, $C_d = (c_1, c_2, \dots, c_m)$ is the set of ontological concepts identified within the document, and $W_d = (w_1, w_2, \dots, w_p)$ is the set of weighs assigned to concepts, where w_t is the relevance weight of the concept c_t for the document D .

Fig. 5 illustrates the processing chain developed to translate XML documents according to this model.

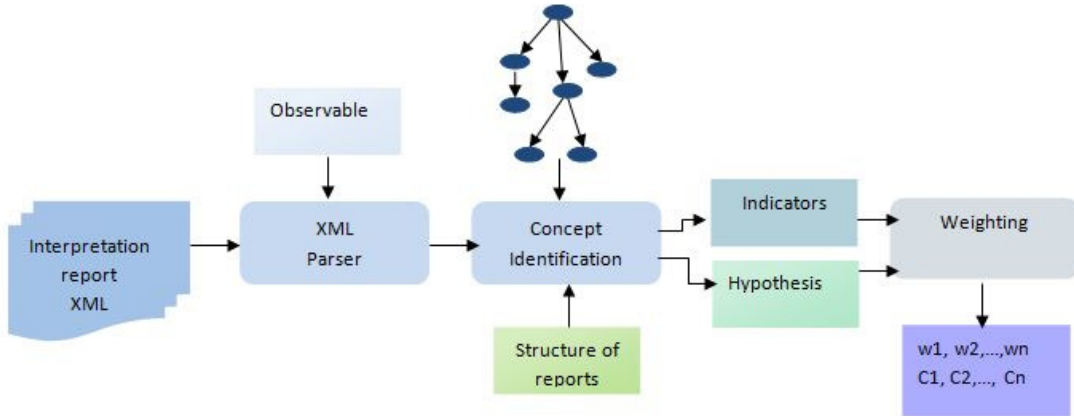


Figure 5: Representation of interpretation reports

First, for each XML file, values of observables are identified by parsing the document. Then, the identification of ontological concepts is carried out by analyzing each semantic unit, and highlighting labels of concepts and relations. This phase identifies two sets of concepts: concepts retrieved within the indicators semantic unit $I = (i_1, i_2, \dots, i_n)$ and concepts retrieved within the hypotheses semantic units $H = (h_1, h_2, \dots, h_n)$.

However, some concepts identified are misplaced as they are not specific to the semantic unit (for example, sub-concepts of *Indicators* retrieved within *Hypotheses* semantic unit),

and they are less relevant. Thus a weighting mechanism is used to assign relevance coefficients to concepts by taking into account the structure of documents and the ontology

Weighting concepts

For each concept set assigned to a document, some concepts appear in their native context, for instance indicators concepts are identified within paragraph describing the conditions of vessel suspicions, while other appear in related contexts, for instance the same indicators related concepts can also appear within paragraphs explaining the link between indicators and hypotheses. We consider that concepts are more relevant when occurred in their native context, and we assign them a numerical value to characterize their degree of relevance. Concepts are automatically weighted as follows:

$$w_t = \frac{1}{|H|} * \frac{depth(t)}{max(depth)} \quad (2)$$

where: w_t is the weight of concept c_t of the set H , $|H|$ is the number of concepts of the set H , $depth(t)$ is the depth of concept c_t in the ontology O , and $max(depth)$ is the maximal depth of the ontology corresponding to its most specific concept. Through this mechanism, the weight 1 will be assigned to concepts identified in their native context (semantic unit), otherwise the weight is estimated by taking into account its specificity with respect to the ontology and its relative importance within the set of concepts. Intuitively, more specific concepts are emphasized, since they carry on a richer semantics. Moreover, their weight is increased when they appear in small-sized concept vectors.

At the end of this phase, is a set of weights associated to the set of concepts identified within the document: $W = (w_1, w_2, \dots, w_n)$.

The goal of those conceptual transition is to represent interpretation reports by a model that keeps the distinction between data and ontological levels, as defined by the index structure.

2.3 Strategies for flexible information retrieval

Two navigation strategies were implemented to explore the collection of documents. Navigation consist of both query definition and information retrieval and it provides improved access to the entire collection by using the index and the representations of previously created documents. The overall navigation process is illustrated in fig. 6.

Navigation starts with a user defining a query expressing its information needs. The query can trigger two information retrieval strategies, according to the set of elements specified by user, and it is executed over the index and the collection of document models. Finally, models relevant with respect to the query are retrieved, and their corresponding documents are presented to users. We adopt a query representation that is similar to document representation model. Hence, a query is a couple

$$Q = (V_q, C_q) \quad (3)$$

where $V_q = (v_1, v_2, \dots, v_n)$ is a vector of observable values, and $C_q = (c_1, c_2, \dots, c_n)$ is a set of concepts. For query processing, there is no weighting of concept relevance according




Figure 6: Strategies for information retrieval

to user interests. Having both the query and documents represented in similar ways, the navigation model implements two information retrieval strategies, called respectively finding and explaining.

Finding facts

Finding in an information retrieval strategy intended to retrieve facts. The input is a user query, and the search strategy is carried out on data level, by exploiting values of observables. The output is a set of documents reporting events occurred in the same region for instance. For this strategy, only the numerical values of the query and the data layer of the index are used. Documents are retrieved by a ranking algorithm evaluating the similarity between the query and models of documents as:

$$Sim(D, Q) = \frac{V_q \cap V_d}{V_q \cup V_d} \quad (4)$$

This strategy is designed for various user profiles and it is suitable when a general overview of incidents is needed. Finding is similar to a browsing task, whereas the user iteratively defines its information need, receives several documents as an answer and uses this new information to refine its information need and improve the navigation.

Explaining behaviours

Explaining corresponds to an investigation search strategy. It receives a user query as input, underlying some behavior considered of interest, and the goal is to identify documents reporting similar behaviors, regardless of their spatial-temporal contexts. For this strategy, both the numerical values of the query and the concepts (in natural language) are used. For this strategy, substantial knowledge is required, and the conceptual layer of the index is used. Documents are retrieved by a ranking algorithm evaluating the similarity of query and document as:

$$Sim(D, Q) = \sum_{i=1}^n w_i * c_i * c_q \quad (5)$$

where n is the number of concepts assigned to document representation w_i is the weight of the concept c_i and $max(e_i, C_q)$ maximizes the semantic similarity between c_i and the set of ontological concepts C_q of the query. The output is a set of documents highlighting various hypotheses to be considered pertinent as to explaining the behavior under analysis. This navigation strategy is designed for particular profiles of users, such as experts investigating maritime incidents, as it could stimulate analogical reasoning. As the data and ontological levels of the index are connected, it becomes possible to perform inferences in order to refine the initial query. Thus, users are able to lead the selection of documents and therefore to expand the set of documents by navigating along the semantic relations. Explaining strategy is similar to a search task that can also be carried out iteratively as returned reports can provide clues to refine the search.

Validation of the model

The evaluation of the retrieval strategies was made in collaboration with maritime security experts and this step was implemented by using jointly the finding and retrieval strategies. The goal of the evaluation was to investigate the impact of using the hybrid index for information retrieval.

The validation methodology followed the following steps:

- Definition of a specific task on illicit trafficking of products;
- Definition of a set of test queries that are suitable for the task. For this step, the user provides a query consisting in a list of numerical values (observable) and keywords for concepts;
- Execution of the set of test queries according to a finding strategy;
- Enrichment of the test queries thanks to the concepts of the ontology. The ontology is used to find relations of those keywords and add them to the initial query.
- Execution of the enriched query according to the explaining strategy;
- Qualitative analysis of the results by experts ;

During the last step, the results of the enriched query were compared to the set of reports retrieved by the initial query. The validation protocol was implemented using the collection of reports created by the project, which is composed of around 37 reports describing 6 types of incidents. Qualitative procedures were needed, given the very small size of the collection.

Evaluating such an approach is a difficult task as we face the users' perception of the results. In this very first series of experiments, we have chosen to focus on the first 15 results returned for each query (top-15).

The analysis of experts showed that enriched queries retrieved more interpretation reports, and this result is in line with the intuition that adding ontological concepts provide a mechanism to retrieve reports that are unreachable through numerical values of variables. The queries that obtained the most different results, when we compare both sets of documents retrieved, are those for which *close to exact* answers were present in the data base of reports. This result is not very surprising since the data layer of our index is a decision tree structure, and sensitivity to outliers [142] is a well-known drawback of decision trees. To conclude, a main advantage of our explaining strategy is the ability to identify common concepts and the joint analysis of numerical values and concepts.

3 Semantic mediation for on-the-fly integration of soft and sensor data

3.1 Dynamic fusion of sensor data and observations

This contribution presents a semantic approach developed to combine soft and sensor data for entity tracking and identification for general security scenarios, such as search and rescue or surveillance. A general architecture was developed for information fusion, which creates a situation using sensor data and enriches this situation by taking into account observations or data provided by humans.

Sources of data

The types of sources available include ground moving target indicator (GMTI), infrared and visible light imaging (IMINT), and signal intelligence (SIGINT) sensors. All of these sensing modalities generate elements that can be described by mathematical and numerical or symbolic representations (e.g., using a universe of discourse), and serve as inputs to automated processing procedures. Sensor measurements result in observations of objects, for which they provide information about properties like location, speed or signal characterization when these objects are electromagnetic emitters. Soft information derived from human or open source is fundamentally different in that its content is often more qualitative and requires additional context elements for complete human interpretation.

Cycles of information fusion

As shown in fig. 7, heterogeneous fusion is carried out by means of two information fusion cycles. The core is a sensor-based kernel that provides several processes for entity correlation and tracking, along with estimation of their states.

The kernel implements a short-time classical tracking algorithm, since data are provided by sensors on a regular frequency. The outcome is a situation, whose entities are described by their spatio-temporal coordinates and their kinematics. At this stage, the type of entities is also estimated but only using sensor-based data. The second layer of this architecture enriches the situation by integrating soft-data elements on a stream and irregular basis, as they become available. The enrichment is aimed at refining the states




Figure 7: Cycles for dynamic information fusion

of entities by adding supplementary attributes, such as allegiance and military or civilian nature. Heterogeneous fusion can be considered as a long-time fusion cycle, triggering specific processes as soft-data observations arrive. Those processes first provide matching mechanisms to assign soft-data observations to entities of the situation and then perform fusion strategies in order to combine elements of entity states with items extracted from soft data.

3.2 Information extraction and semantic mediation

Extraction of features from sensor data

The features that can be extracted from sensor data, although limited, are heterogeneous due to the different types of deployed sensors. GMTI sensors are radars with specific signal processing leveraging the Doppler effect which can provide information about moving targets, mainly related to their kinematic state (location and speed in the direction of the sensor) and sometimes some classification information from signal analysis, limited to rough classes of identification, such as rotating objects (e.g., helicopter blades), tracked vehicles such as Tanks, or wheeled vehicles. The location is provided in range and azimuth and possibly in elevation for a 3-D radar, with associated imprecision in each of these dimensions, which can be quite large for azimuth information. The speed is also partially retrieved due to the fact that only one-dimensional information can be acquired, related to the line of sight between the radar and the detected object.

IMINT information is related to imagery or video acquired in diverse wavelengths (visible or infrared), which may be further exploited through an automated extraction and tracking device and annotated by an operational user. Some sensors are also able to perform tracking on a given object, thus providing track information about the object with location and speed attributes. The operator can then complement this information

with precise classification information.

SIGINT information is related to either the detection and localization of specific emitters (e.g., radars) or the interception of communication information between several actors. Through the use of several receivers or a maneuvering receiver, the location of emitters and specific technical information can be extracted. From the set of items sent by these various types of sensors, a correlation scheme can be set up based mainly upon kinematic information to perform the right association between these detections, so there are a limited number of duplicated tracks related to the same real entity. From this association, a combined estimation of attributes is performed, which leads to a more precise kinematic information and, if available, a rough estimate of the type and hostility of the entity. Different methods, using mainly Bayesian, evidence theory or heuristic techniques, are involved to combine classification results from these detections.

Identification of properties from soft data

Processing of soft data identifies properties of entities within natural language messages. Messages also have a heterogeneous content, and can provide insight on different aspects, such as entity location, and evolution. The methods developed extract binary association attribute-value from messages, which can be easily modeled as properties of entities and integrated into entity states. Attributes specify the type (vehicle, bus, person, etc.), allegiance (foe, friend, neutral) or nature (civilian, military, insurgent, etc.) of entities. Attributes are identified from soft data by using a text-mining approach based on collocation identification. Collocations are associations of words that co-occur frequently within the same sentence, whether because the meanings of words are related to each other (e.g., vehicle- road, car-driver) or because the two words make up a compound noun (car stop, subway station). The extraction algorithms focus on collocations composed of two words, also called bi-grams. The method developed to extract collocation is shown in fig. 8.




Figure 8: Extraction of collocations

Because simply taking the entire list of collocations captures an excess of extraneous and incoherent information, additional processing is needed to filter relevant word associations thanks to semantic annotation. Semantic annotation is performed automatically, using procedures based on lexical similarities, which associate a real number with a pair of words. Lexical similarity offers a measure of the degree to which two words are similar and are used to label a collocation by ontological concepts. The hierarchy of concepts allows the attribute part of the association to be retrieved. Thus, given a concept C considered as value, the associated attribute is identified as the least specific concept subsuming $Property$ and generalizing C as shown in fig. 9.

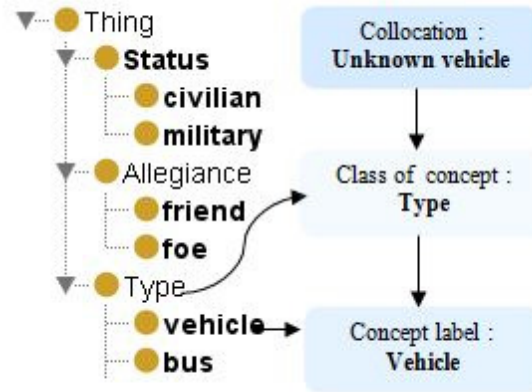


Figure 9: Semantic annotation

For instance, the collocation unknown bus will be matched to the set $(Bus, Bus, Type)$, since unknown is not assigned to a concept, while the insurgent vehicle is annotated by $(Vehicle, Vehicle, Type)$ and $(Insurgent, Insurgent, Allegiance)$, since both Vehicle and Insurgent are concepts of the ontology.

At the end of this phase, annotations of collocations are generated in the form of tuples: $A_i = (W_i, C_i, T_i)$ where A_i is the annotation of item i , W_i is a word, part of a collocation, C_i , is a concept assigned to W_i by lexical similarities, and T_i is the category of C_i , as identified by inferences

Integration of sensor data and observations

The overall solution for human observations and sensor data integration consists in the assignment of observations sent by humans to entities created by sensor data processing. Soft data are brief human reports, conveying information about entities in the field. Among those entities, some could be of interest, already detected and eventually tracked by sensors. Information extracted from incoming soft reports will not be considered for tracking purposes; instead, it will help human-operators to have a better description of the situation.

Assignment of human observations to entities of the situation is carried out in the light of spatio-temporal correlation. Given that observations are associated with a timestamp and have specific locations, this method first estimates a correlation coefficient to describe the probability of a human observation to be assigned to e_i , an entity of the situation. The current states of e_i along with its previous states are taken into account for this estimation, since soft observations are not necessarily synchronized with the current situation. Results of this estimation are then ranked and the observation maximizing the value is selected and added to the set of observations associated with e_i .

Semantic mediation

Semantic mediation relies on using a domain ontology to describe data semantics, and implementing semantic annotation procedures to associate sensor output and human observations with corresponding elements of the domain ontology along with reasoning mechanisms for data integration. While spatio-temporal association enriches the overall situation by adding a set of human observations, semantic mediation is used at entity level to enable the fusion of sensor data and observations. Semantic mediation is implemented as a process allowing data provided by different types of sources to be combined and a domain ontology is at core of this process, as shown in fig. 10.

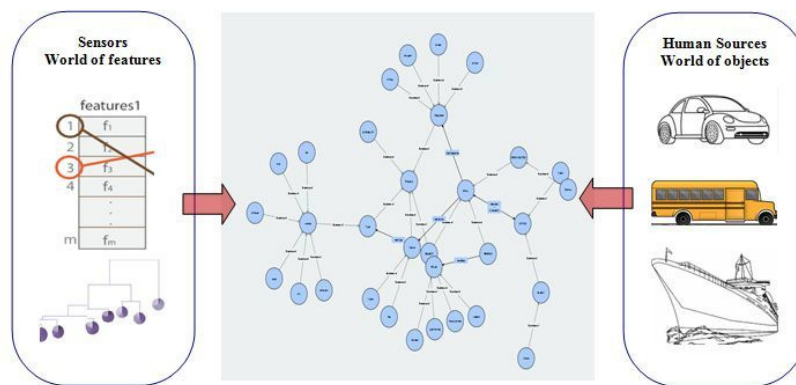


Figure 10: semantic mediation from the *world of features* to the *world of objects*

In the context of this work, the role of the mediation process is to integrate relevant observations to sensor inferred entities according to a shared semantics modeled by a domain ontology and to master the gap between low-level features and richer conceptual descriptions of each entity.

The ontology used for this work was created by using a top-down approach. The development began with a preliminary conceptualization, where a list of high-level concepts was identified by using the MIM Information Model ⁷. The hierarchy of concepts was iteratively enriched by adding new classes. The result is a domain ontology composed of 31 concepts, see fig. 11.

Main classes of the ontology are *Entity*, *Action* and *Property*. The ontology also has 6 object properties, to make explicit interactions between classes. Thus, $hasActor(Event, Entity)$ models the association of an *Entity* with an *Event* and its reversed relation is $involvedIn(Entity, Event)$. The four remaining relations associate entities and their properties: $hasRole(Entity, Role)$, $hasType(Entity, Type)$, $hasNature(Entity, Nature)$ and $hasCapability(Entity, Capability)$. All developments and testing were carried out using Protégé, and the ontology is represented in OWL DL, a description logic [1] sub-language of OWL [6].

⁷<https://www.mimworld.org/portal/projects/welcome/wiki/Welcome>

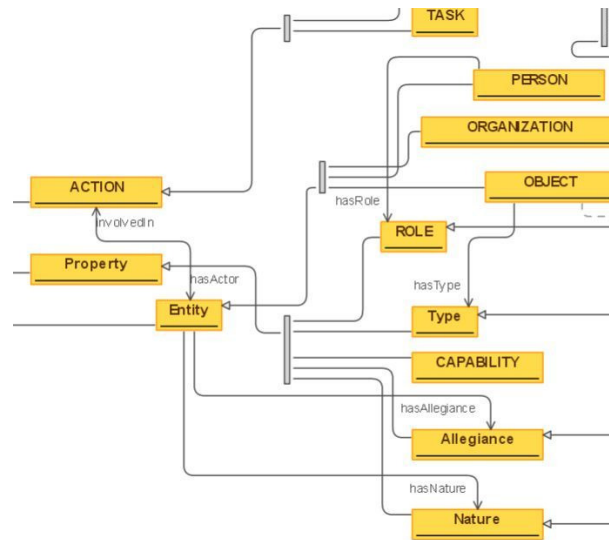


Figure 11: Ontology for semantic mediation

Refinement of entity states

In the surveillance problem described in this contribution, complementary and overlapping inputs exist. Sensor processing has the capacity to identify the type of entities, which is then added to the entity state in the form of a tuple attribute, value-of-attribute, for instance type, bus.

Sensors classify entities based on their measured features and by using some supervised methods for classification, which provide a limited number of categories. Besides having a limited number of categories, sensor processing also has its own detection limitations, and more subtle aspects such as the allegiance of entities are out of reach for their sensing capabilities.

Fusion of sensor and soft data is twofold: first, complementary properties extracted from human reports are added to states of entities; second, the type of entities is updated by taking into account the type of entities as stated by sensor processing and the type as extracted from soft data. In order to update the type, reasoning mechanisms are used to combine attributes of entities. More specifically, given that operators are interested in having a more precise description of entities, reasoning procedures identify the most specific concept of both type labels. This concept is then used to describe the entity, as illustrated in fig. 12, where the final state of the entity highlights the type *bus*, as a concept more specific than *vehicle*.

Observations can be noisy, incomplete and sometimes irrelevant, and the inference mechanism fails to identify the most specific concept of both type labels. In this case, fusion provides inconclusive results. Inconclusive inference is due to contradictions between sensor reports and observations, or accidental associations of observations with entities. When successful, the result of fusing soft and sensor data at entity level is a more specific identification of the type of entities, and the enrichment of their state thanks to additional properties that cannot be inferred by sensor processing.




Figure 12: Refinement of entity states

3.3 The INFOTEX prototype

The approach of dynamic information fusion was implemented by using FUTHANE, a generic development and execution framework providing a collection of basic algorithmic building-blocks for information fusion. The result of this implementation was the INFO TEX prototype.

An experimental track-and-detect scenario was adopted to provide a valid proof-of-concept of the fusion system. The scenario involves a multi-sensor multi-tracking task with a network of sensors, several observers and main capabilities for information fusion. A total of 20 observations sent by operators on the ground were used, in addition to 204 GMTI reports, 12 COMINT reports, 7 ELINT reports and 5 IMINT reports. At entity level, properties extracted from soft data describe allegiance (*friend, foe, insurgent*), type and nature (*civilian, military*), as shown in fig. 13.

AdditionalProperties	
Property	
Name	Allegiance
Value	insurgent
Property	
Name	Type
Value	ATV
Property	
Name	Nature
Value	civilian

Figure 13: Attributes of entities for situation awareness

Sensor data is provided as formatted reports integrating position, time, type and potential subtype of an observed entity. Some additional features may also be present, such as vehicle color or even the vehicle identification number. Observations are in the form of structured reports, having a natural language paragraph to summarize information collected by human sources. After feature extraction from text and fusion of items, the

type of entity is updated (Tank) and its hostility is identified (HO). Over time, 102 entities are detected and tracked in the situation.

Although using an experimental scenario offers a basis to evaluate whether or not the system meets its objectives, a formal evaluation is needed to estimate the impact of using semantic mediation to support heterogeneous fusion. Since mediation affects both entity states and the overall situation, information gain and quality of service introduced in [126] to characterize situation assessment are metrics able to quantify this impact.

Information gain is a criterion intended to capture the value added to entity states after updating their descriptions by using semantic inferences. Information gain is the ability of the system to provide improvements, and its values can be assessed by taking into account the number of additional properties added to entities and the quality of their type. *Information gain* is defined assuming that changing the state of entities by integrating observations improves the description of entities, as follows:

$$InfoGain = \frac{N_c}{N} \quad (6)$$

where N_c is the number of entities whose states are affected by observations, and N is the overall number of entities of the situation. Values of *InfoGain* range from 0, when observations are not related to entities of the situation, and 1, when ideally states of all entities are updated.

Quality of service is a criterion used to characterize the quality of situation assessment and encompasses aspects related to timeliness, uncertainty of the overall picture, and quality of individual descriptions at entity level. Nevertheless, uncertainty can also arise at the situation level, since soft and sensor data can provide contradictory information items. Quality of service takes into account the number of failures due to inconclusive fusion:

$$QoS = \frac{U}{N_c} \quad (7)$$

where N_c is the number of entities whose states are affected by observations, and U is the number of valid inferences. *QoS* ranges between 0, when all inferences for type refinement are inconclusive, and 1, when they are valid. The evaluation scenario was run by iteratively including 20 observations, and the evolution of those criteria is shown in fig. 14.

The values of information gain and quality of service at the end of the scenario are 0.47 and 1, respectively. For this experimentation, all inferences for type refining are valid and low values of information gain are directly related to the small number of entities affected by incoming observations.

4 Related work

Using semantics to deal with heterogeneity and dynamic aspects in complex environments is an active research field, and several ontologies have been already created as a result. Among them, SSN is an ontology created by the Semantic Sensor Networks Incubator




Figure 14: Evolution of Information gain and Quality of service

Group ⁸ offering descriptions of four related perspectives: sensor, observations, system and property [13]. The model focuses on what and how can be sensed, systems of sensors and their deployment, and the description of observations made at entity level or for particular properties.

Sensor Web Enablement (SWE) is a major initiative of the Open Geospatial Consortium (OGC) ⁹ aimed at improving the capacity to discover relevant sensor data on the Web, through standardized interfaces and specifications, and creating the Semantic Sensor Web, an infrastructure that enables the interoperable usage of sensors by providing services for discovery, accessing and identification based on sensor output augmented with spatial, temporal, and topic-specific semantic metadata [49]. A complement of the Semantic Sensor Web, Sensor Linked Data is a paradigm introduced by Janowicz and colleagues [109] to add interlink as a new challenge. The SOSA ontology [17] is the latest model, developed up to date, to link observations, sensors and measures.

More related to the second approach presented in this chapter, various research efforts consider the existing bridge between symbolic knowledge representations and raw data collected and measured by sensors. The concept of semantic perception is at the core of approaches developed to upgrade sensor output by attaching semantics to sensor data for enhanced interpretation [119]. Roda and Musulin present in [102] an ontology-based framework to support intelligent data analysis of sensed data, while the construction of a complex situation using multi-layer ontologies is discussed in [60]. Semantic mediation in sensor networks is addressed by Malewsky [112], who developed a semantic framework to improve matchmaking on sensor measurement and operating capabilities and used a modularized Sensor Mediation ontology aligned to the SSN.

The rise of the Internet of Things (IoT) offers a wide application area for fusion methods able to discover potential knowledge from large amounts of perceptual information [67]. Within this context, semantic-based solutions are used to add semantics to data [43] or to support event matching [99].

⁸<https://www.w3.org/2005/Incubator/ssn/>

⁹<http://www.opengeospatial.org/>

Internally, within the frame of the FUTHANE project, another research line investigated the use of ontologies to filter relations extracted from textual messages, in an effort to add not only entities in the situation but also relations describing their interactions. The approach is described in [69], while experiments and results are discussed in [71].

5 Lessons learnt

In this chapter I drew up an overview of two solutions using semantic interoperability in order to improve the analysis of heterogeneous and dynamic situations. With the example of information retrieval for the maritime surveillance, we have shown how semantic processing can be employed to build more intelligent integrated systems.

Developing the semantic interoperability approach for maritime surveillance domain was my main contribution to the I2C European project. I came with the idea of connecting indexing and retrieval strategies by interviewing several domain experts, and in particular maritime operators working with DCNS (currently Naval Group)¹⁰, the partner leading the project. For domain experts, the problem encountered was to connect the storage and retrieval of information together and in practice this need was satisfied by adding (lots of) data formatting and data base query steps, making the integrated analysis of items difficult to achieve.

Designing and implementing the semantic mediation system for soft and sensor integration was the main result of InfoTex, an advanced study of the FUTHANE research program, funded by the French MoD. For this specific application, the use of semantic integration lead to the development of new capabilities for information fusion, that were greatly appreciated during the evaluation of the INFOTEX prototype by end-users.

Special attention was given to the evaluation task and, as the scientific coordinator of the FUTHANE project, I struggled to keep the distinction between the teams that developed and assessed the system, respectively. In this respect, I supervised Thibaut Castaing, a post-doctoral student who joined the project for 9 months and was assigned specifically to the evaluation of the INFOTEX prototype.

Both applications show that incorporating domain knowledge makes the processing task more complex and thus more time consuming.

6 Collaborations and dissemination of results

I2C European Project, FP7

Semantic integration for the maritime domain was a task of the I2C European project. I collaborated with experts from Naval group and also with my colleagues from the DEMR department (Département électromagnétisme et radar) [98].

¹⁰<https://www.naval-group.com/fr/>

FUTHANE, Advanced Research Project, French MoD

I was the scientific coordinator of the FUTHANE project and the principal investigator of InfoTex. I coordinated and I have actively worked with the mixt ONERA - AIRBUS team in charge of the InfoTex study, on specifying the fusion algorithm and carefully defining the structure of the mediation framework. I had close collaborations with **Amandine Bellenger**, **Frederic Gallésio**, **Xavier Lerouvreur** and **Sylvain Gatepaille**, [84], [51], [86].

Supervision of students

I also supervised **Thibault Castaing** during his 9 months post doc within the frame of the FUTHANE project.

Dissemination of result

Dissemination of activities presented in this chapter include:

- A book chapter in *Shahbazian, E., and Rogova, G. (Eds.). (2016). Meeting Security Challenges Through Data Analytics and Decision Support (Vol. 47). IOS Press ;*
- A journal article published in the *International Journal of Knowledge and Systems Science*, [70];
- Demonstrations and presentation of results to the French MoD, during the assessment of the FUTHANE project;
- Presentation of the semantic mediation solution in the context of the *Workshop on AI and Cognitive Technologies for Radar, Communications and and EW* organised by the European Defence Agency (EDA), Brussels, December 2018, [34]

Main publications of the chapter

- [51] Valentina Dragos and Sylvain Gatepaille. “On-the-fly integration of soft and sensor data for enhanced situation assessment”. In: *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 21st International Conference KES-2017*. Vol. 112. Procedia Computer Science. Elsevier, 2017, pp. 1263–1272.
- [69] Valentina Dragos. “Comparative Studies on Using Semantic Filtering for Open Relation Identification.” In: *Meeting Security Challenges Through Data Analytics and Decision Support*. IOS Press, 2016, pp. 69–77.
- [71] Valentina Dragos, Sylvain Gatepaille, and Xavier Lerouvreur. “Refining relation identification by combining soft and sensor data”. In: *2016 19th International Conference on Information Fusion (FUSION)*. IEEE. 2016, pp. 2139–2146.

- [84] Valentina Dragos, Xavier Lerouvreur, and Sylvain Gatepaille. “A critical assessment of two methods for heterogeneous information fusion”. In: *Information Fusion (Fusion), 2015 18th International Conference on*. IEEE. 2015, pp. 42–49.
- [86] Sylvain Gatepaille, Valentina Dragos, and Frédéric Dambreville. “FUTHANE: An Open Framework to Support Heterogeneous Information Fusion”. In: *2015 NATO IST-SET-126 Symposium on "Information Fusion (Hard and Soft) for ISR*. IEEE, Norfolk Virginia, United States. 2015, pp. 21–30.

Chapter 4

Content characterisation for cyberspace exploration

Contents

1	Motivations	57
2	Formal models for online content characterization	59
2.1	An ontology of appraisals for subjectivity modeling	59
2.2	Conceptual foundations of hate modeling in French	64
3	Illustration of ontology-driven social data analysis	69
3.1	Detection of trends in online data	69
4	Related work	75
5	Lessons learnt	76
6	Collaborations and dissemination of results	77

1 Motivations

This chapter addresses the exploration of the cyber-social space as defined by Shets and colleagues in [114], as *a network of humans and autonomous agents and theirs links creating human, autonomous or hybrid communities*. Recent examples on disinformation [45] or social manipulation [19] show that phenomena in the cyberspace have the capacity to polarize social views, making social groups form and fracture in online spaces with concrete consequences on political and social environments.

A critical aspect of social media misuse is the spread of content related to extremist attitudes and ideologies, including online propaganda, hate speech or radicalization. Often, this content involves the presence of opinions, beliefs and other subjective states that play a key role in the sentence. While harmful content is removed every now and then on social platforms and Twitter actively shuts down extremist handles, the account removal approach is not working despite its short-term effectiveness and claims of long-term efficacy [82]. From a defence and security perspective, it is then important to have means to

effectively mine the cyberspace for features of emotions, opinions or sentiments [74], [53]. Approaches focusing not on *who is posting*, but rather on *what is posted* are needed.

One of the major promises for users when using the cyberspace as an information source is that data are easier to collect, share and reuse. However, on concrete situations, social data do not appear to be effectively processed and exploited as much as they could be, mainly because such data poses to the community the challenge of *building reliable systems on unreliable data* [93]. A large part of my work in the last years has been dedicated to the topic of social data analysis, and more precisely I tried to understand (i) how (and how much) do online users provide useful information on social platforms, (ii) why end users may not be able to exploit social data as much (and as frequently) as they wanted to and (iii) what can be done to enhance the analysis of social data for defence and security applications.

Most current methods and techniques developed for social data analysis rely strongly upon linguistic inputs, making use of blacklists, regular expressions, annotated corpora or lexicons. Those techniques have specific performances and drawbacks: while keyword can be indicators for hateful comments, alone they are not enough to detect propaganda, and often dictionary-based approaches identify accurately the target of speeches but fail to distinguish sentences conveying extremist messages from the clean ones.

From a different perspective, artificial intelligence techniques relying on statistical methods based on co-occurrences of words or purely linguistic approaches are unable to deal with this type of content, and are far from being able to infer the subjective and cognitive information associated with social data. For instance, *religious-belief* and *buy-Easter-chocolate* are two expressions with subjective connotations that come not from individual terms, but from putting those terms together. Techniques that rely on mere annotation would perform poorly in discerning the meaning of such expression, and their effective processing requires a significant amount of knowledge.

The originality of my approaches lies in closely working with researchers on sociology and linguistics, in an effort to capture clues that help the characterization of online content and developing semantics-based techniques to address the problems identified. Several semantic resources have been developed to help the analysis of social data in English and French. Although the research trend is dynamic and several resources are available for subjectivity analysis in English [42], [125], it should be noted that there are very few corpora and semantic resources for French in the more general area of social data analysis [46], [135], [132]. Another contribution of my work was to develop resources to analyze online contents in French.

This chapter introduces several contributions. The first section presents two ontologies: the first one was developed to capture features of subjectivity and the second one is intended to support the analysis of online hate. The ontology of hate was developed specifically for French corpora. Pitfalls and limitations of building ontologies for subjectivity analysis are also discussed for both models. The second section of this chapter presents an illustration using ontologies and unsupervised techniques to detect trends in social data.

2 Formal models for online content characterization

2.1 An ontology of appraisals for subjectivity modeling

This contribution describes the construction of the appraisal ontology, a formal model grounded on cognitive foundations that offers a mean to detect subjective stances and harness affective states within content released by users on the Internet. The ontology is built upon the foundations of the appraisal theory [94] developed by White and Martin. The appraisal frame considers the meanings by which text convey positive and negative attitudes, and takes into account the way such utterances are strengthened or weakened by author's thanks to linguistic clues of intensity. Adopting appraisal categories allows going beyond the limitations of opinion and sentiment, and presents author's feelings, tastes and opinions with greater or lesser intensity.

Cognitive foundations of the model: the appraisal theory

The appraisal theory is a cognitive frame claiming that people's emotions are elicited by their personal and continuous interpretations, evaluations or appraisals of objects, events and situations. The most important aspect of the appraisal theory is that it provides a way to express how humans interpret some particular event as positive or negative along with their position, support and engagement with respect to their own interpretation and report: confidence, support, agreement, disagreement, deny, rejection. From a linguistic standpoint, the appraisal theory describes how authors use linguistic expressions to communicate their emotional states and engagement. The appraisal framework goes beyond limitations of traditional concepts of sentiment and opinions and considers the appraisal expression, which is a linguistic unit by which a personal appreciation is conveyed, whether it is an opinion, sentiment or supportive statement. The appraisal theory structures appraisal expressions under three main basic systems describing attitudes, engagement and graduation, as shown in fig. 1.




Figure 1: Systems of the appraisal theory

The *attitude* system is related to linguistic expressions conveying the current of authors at the time they write the text and. This system covers three main subcategories: *affect*,

appreciation and *judgement*, discussed hereafter.

Affect is related to linguistic expressions of author's feelings such as happiness, joy, sadness, grief, etc., as shown in sentences hereafter:

*S1: This movie **bores** me.*

*S2: Helping others makes me **happy**.*

Judgment highlights linguistic expressions conveying characterization of persons and behaviors by the author. Generally it conveys opinions and personal tastes about objects, such as nice, ugly, beautiful, shy but also about interactions and behaviors in the social context: heroic, brave, open-minded, feeble minded, see sentences below:

*S3: Paul is **rude**, but **honest**.*

*S4: They are a **great** nation.*

Appreciation is related to assessment and evaluations of entities, objects, events and scenes, as shown in examples:

*S5: The painting is **beautiful**.*

*S6: My phone is **useless**.*

The *engagement* system gathers linguistic expression specifying the author's position with respect to his own statements see fig. 2. When reporting, writers often embed clues as to how strongly they support the content being conveyed and may indicate confidence, doubt, skepticism, conviction, etc., about the information reported.

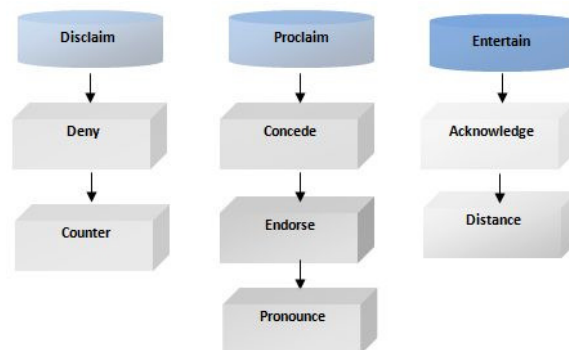


Figure 2: The Engagement System

The engagement system is closely related to the notions of trust, confidence, probability or possibility. Categories under this system encompass aspects related to denial (S7), concession (S8), confirmation (S9), endorsement (S10), acknowledgement (S11) and distance (S12).

*S 7: You **don't** need to access the file.*

*S8: **Although** it was raining we went out.*

*S9: **Of course** Brexit was unexpected.*

*S10: The reports clearly **show** he was involved in the accident.*

*S11: **According to** Times, he is leading the election.*

*S12: **Many** are claiming that he will not win.*

Graduation is the last system of the appraisal theory and it is introduced in order to provide means to measure or at least estimate the orientation and various degrees

of intensity associated to affect, appreciations, judgment and engagement. The system introduces concepts to describe the intensity of opinions. Main concepts are *Force* and *Focus*, and the graduation system also models linguistic modifiers such as intensifiers (very, enough, etc.) and downtoners (few, low, etc.).

A formal model of appraisal categories

The appraisal ontology was created in order to build a formal model of appraisal categories. The ontology is a commonsense knowledge base intended for subjectivity analysis that specifies the systems of the appraisal theory.

We first created an initial version of the ontology by taking into account the appraisal theory. We started with appraisal categories and then added additional concepts to characterise appraisal expressions. As a result, six main concepts were identified and modeled: attitude, engagement, graduation, orientation, polarity and modifier, as shown in fig. 3.



Figure 3: The Appraisal ontology: main concepts

Among appraisal categories, attitude and engagement are modeled as introduced by the appraisal theory, but with finer descriptions. Sub-classes of Attitude and Engagement systems are shown in tab. 1 and tab. 2, respectively.

Classes	SubClasses	Instances
Appreciation	Composition	Detailed, consistent
	Impact	Amazing, monotonous
	Valuation	Innovative, profound
Judgement	Social esteem	Brave, disloyal , clever
	Social sanction	Virtuous, corrupt, honest

Table 1: Refining classes of the Attitude system

Classes	SubClasses	Instances
Disclaim	Deny Counter	No, didn't, never Yet, although, but
Proclaim	Concur Pronounce Endorse	Naturally, obviously Indeed Demonstrates, shows, proves
Refer	Aknowledge Distance	It's said that, many argues He claimed to

Table 2: Refining classes of the Engagement system

Graduation concept was enriched by adding several subclasses to describe *Focus* as *Sharp* or *Soft* and *Force* as *High* or *Low*. For all those categories linguistic examples were also added at instance level. *Modifier* is a concept intended to capture the ability of linguistic markers to increase, decrease or reverse the intensity of appraisal expressions and *Polarity* is a concept used to capture positive and negative emotional states conveyed by appraisal expressions.

Besides *is-a* relations used to create the hierarchical structure of concepts, 6 relations were added to connect concepts on the ontology, see tab. 3.

Classes	SubClasses	Instances
Name	Source	Target
hasGraduation	Engagement	Graduation
hasPolarity	Engagement	Polarity
hasOrientation	Engagement	Orientation
Increase	Modifier	Force
Decrease	Modifier	Force

Table 3: The appraisal ontology: relationships

By considering only the appraisal theory, the appraisal ontology is composed of 46 concepts with a 6 levels hierarchy; the model also has 4 Object Properties and 2 DataType Propertie, 50 nodes, 75 edges and 268 instances of concepts. The polarity, force and focus of a word were added manually to concepts of the ontology.

Increasing the robustness with WordNet

We observed that this version, although being in line with the appraisal theory, has a main weakness: a word can have synonyms, and choosing the set of words covering a meaning is important for the purpose of increasing the accuracy of the appraisal ontology. To overcome this issues, we improved the robustness of the ontology by using WordNet [157].

WordNet is a lexical data base for English, where each word is associated with one or more synsets. A synset consists of a identifier that represents a unique meaning. A word

can have one or more synsets (love can be in reference to the verb or the noun and thus related to two different synsets) and one synset can be related to one or several words (nice and pretty can be synonymous). The enrichment followed three main steps:

- Select the first WordNet entry for a concept of the ontology;
- Select the most semantically related synsets of this entry . To estimate the similarity of synsets we used the Leacock Chodorov (LCH) similarity [156];
- Select instances to be added to the ontology from the synsets previously retrieved.

Pitfalls and limitations of the appraisal ontology

Practical pitfalls affected the construction of the appraisal ontology and the model has several limitations, discussed hereafter.

- Modeling nonfunctional domains

The ontology was constructed to capture features of subjectivity according to principles of the appraisal theory. However, the model is not build to describe a functional domain, such as aeronautics or space, having well defined concepts and interactions. Thus, the ontology has a limited number of horizontal relations, describing interactions of concepts and strives to represent the terminological components of concepts. The terminological dimension is emphasized by adding instances of concepts, and also by highlighting how a broader or narrower meaning can be created by linguistic adjustments. Moreover, elements of the model are drawn from the realm of thesauri and semantic lexicons, as there are no sources to describe subjectivity and there are no domain experts to guide the modeling process.

- Impact of linguistic phenomena

The main impact from linguistic standpoint is due to the fact that the meaning of words usually depends on the context in which they occur. More specifically, the polarity of words can change according to the context in which they are being used. Although polarity of *good*, *bad* or *excellent* remains the same in all context, *small* is a example of polarity ambiguous word, having a positive orientation in sentences like *small desynchronization*, rather negative when used as *small capacity* and rather neutral in *small world*. As polarity is modeled as a concept of the ontology having three exclusive values (positive, negative and neutral), overcoming limitations dues to ambiguous polarity requires the implementation external procedures, able to detect the occurrence context of words and infer their polarity accordingly.

- Completeness and relevance of the model

The final set of limitations concerns the development of a computational model of the appraisal theory. First, there are limitations intrinsic to the cognitive framework, according to which events can trigger more than one affective response simultaneously

with various intensity levels although the theory does not provide a clear strategy to select an affective state during the appraisal process. This limitation has direct impacts on the modeling process, as there is no guidance as to how to add relations in the ontology, how to add instances of concepts and how to craft a set of rules for inferences. To overcome this limitation, a solution providing ethical procedures to guide the practical use of the model is discussed in [120]. In addition to limitations cited above, the use of ontology for social data analysis might be able to also highlight task-related drawbacks, such as modeling only concepts and relations in English, although some social media messages are a mix of words written in several languages.

2.2 Conceptual foundations of hate modeling in French

The ontology of appraisals introduced in the previous contribution allows the analysis of subjective states in online data. This contribution takes a step further and discusses the construction of a model that will support the detection of online hate. Detection of hate is part of risk prevention strategies for defence and security practitioners. For those domains, there is a practical need to develop semi-automatic approaches for online hate investigation.

Resources and methodology

The construction methodology consists of three main phases: acquisition of knowledge, construction of a conceptual model and its consolidation, see fig. 4.

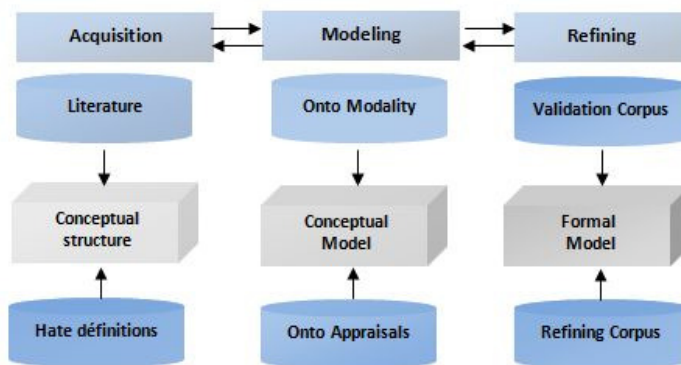


Figure 4: Methodology and resources

During the acquisition phase, three definitions for hate speech were selected in order to gain better understanding of hate semantics in different communities.

The first definition is proposed by Nockleby [154], who defines hate speech as an *abusive speech targeting specific group characteristics, such as ethnic origin, religion, gender or sexual orientation*.

The second definition is given by Mondal et al. who define hate speech as *an offensive post, motivated, in whole or in a part, by the writer's bias against an aspect of a group of people* [59].

The third definition selected for this work is developed jointly by Facebook, Twitter and Google; an online hate speech is defined as *a type of speech that takes place online, generally social media or the Internet, with the purpose of attacking a person or a group on the basis of attributes such as race, religion, ethnic origin, sexual orientation, disability, or gender.*

In the light of definitions above, we consider hate speech as a combination of four main elements : an enunciator - or author - of a hateful speech, a hateful action, a target and the linguistic context of validation of the hateful action. The result of this acquisition phase is a set of main concepts called conceptual structure. During the representation phase, this initial structure is enriched by exploiting two existing ontologies.

The ontology of modalities [131] was developed in order to assess the validity of the knowledge extracted from textual content in the scientific literature but is not specific to this kind of textual data. The ontology describes *the linguistic context of validation* of an item of information, see fig. 5. More specifically, the model represents the relations between three linguistic categories involved in the characterization of validity conditions of information: enunciative conditions, aspect and temporal conditions and also modal (or rhetorical) conditions.

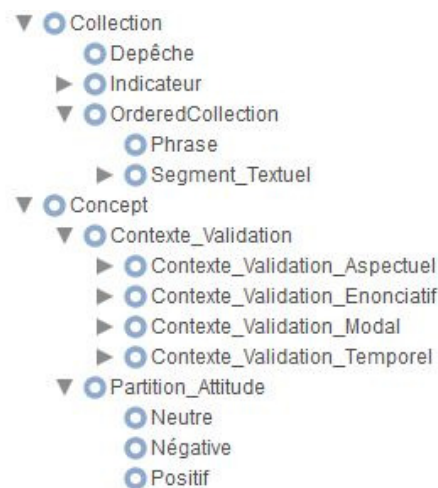


Figure 5: Ontology of modalities

This ontology of modalities models several concepts that are relevant for the purposes of this work. First, the root concept of this ontology, *ContextOValidation* and his two sub concepts, the *LinguisticContextOfValidation* and the *ExtraLinguisticContextOfValidation*, respectively. Moreover, the *LinguisticContextofValidation* deals with the knowledge that is relative to the context of validation as expressed by the author of the text, whereas the *ExtraLinguisticContextofValidation* looks at external knowledge or resources.

Taking into account the enunciative aspect enables us to identify whether the author is fully committed to items in the sentences, or whether he just provides items attributed to another enunciator, with a marker of agreement or disagreement. The epistemic and temporal aspects captured by the ontology of modalities, provide additional attributes of the validation context, for example by capturing whether actions are temporally in the

past time or in the present time, or by capturing degrees of certainty or uncertainty of actions from the point of view of the enunciator.

The aspects captured by the ontology of modalities are relevant for any domain. Thus this ontology is independent of any domain ontology and reusable for other applications. The model has been also used for analysing newswire texts as described in [105]. The resource was built by focusing on the explicit linguistic markers of different kinds of epistemic qualification and it captures temporal, modal and enunciative and aspect - related features which are indicative of authorial commitment to the information conveyed by sentences.

The ontology of modalities and the appraisal ontology were jointly used during the modeling phase to build a novel ontology named 'ontology of hate'. The intuition behind the joint utilization of those resources is that, while the appraisal ontology can provide a set of concepts related to affect or judgement that are of interest for hate modeling, the modalities ontology can facilitate the analysis of their linguistic context of validation, thanks mainly to concepts capturing modality and enunciativity features. Our assumption is that the analysis of this linguistic context helps recognizing reported speech, counter speech, negation, linguistic reappropriation, as well as several modal processes a locutor could use to distance himself from the statement, or limitate its strength.

From these two ontologies, only the relevant concepts for hate speech detection have been kept. At the end of this modeling phase, concepts are added to initial structure and a formal model is built.

The last phase refines the formal model by using two distinct corpora collected on social media. A first corpus was used to manually refine the concepts of the ontology, and this corpora was provided by external contributors. For the validation step, a corpus of tweets was collected using the Twitter API with a combination of relevant keywords. This corpus was collected based on query terms, without manual processing or tagging. The data set is composed of around 300 tweets in French that have been manually selected for their compatibility with hate speech definitions.

Development and description of the model

The ontology was created from scratch, starting with concepts selected from the modalities and appraisal ontologies, and adding additional concepts to characterize online hate. The ontology of hate has four main concepts, see 6: Action (Action), Target (Cible), Context (Contexte) and Orientation (Orientation). The enunciator is part of the *Context* concept.

Action is the core concept to detect hateful content. It is further divided into instigation to violence (AppelALaViolence) or deprecatory judgement (JugementDévalorisant). This last concept is then composed of several sub-concepts extracted from the appraisal ontology. Affect (Affect) captures emotions of the author while Judgement (Jugement) describes assessments of persons and objects and is divided into: social esteem (Estime-Sociale), which encompasses capacities (Capacité) e.g stupid, tenacity (Ténacité) e.g unreliable and normality (Normalité) e.g unpopular, and social sanction (SanctionSociale) which captures truthfulness (Véracité) e.g liar and propriety (Bienseance) e.g corrupt. Another sub-concept of *Action* is *Insult*, a concept regrouping insults which do not involve a specific evaluation of the insulted object but rather derogatory terms for groups



Figure 6: Main concepts of the ontology

of individuals, as opposed to appraisal-oriented insults, e.g jackass (imbécile), carrying specific judgment.

Target (Cible) classifies the targets of hateful statements by taking into account the following characteristics: Sexual Orientation, Ethnical Origin, National Origin, Sex, Religion, and Immigration Status.

Context is a larger umbrella for several concepts extracted from the ontology of modalities. It describes: the aspectual context (*ContexteValidationAspectuel*) of *Action* given by the author and which can be for example a process, an event or a state ; the enunciating context (*ContexteValidationEnonciatif*) which can be *Collectif (collective)* in cases of factual statements and (so called) global or wide spread opinions or *Individuel (individual)* when the author or another source cited by the author is implied in a hateful action; the temporal context (*ContexteValidationTemporel*), capturing the absolute or relative temporal of statements which can be in the past, the present or the future of the enunciator; the modal context (*ContexteValidationModal*), which carries the author’s engagements towards its statements.

Moreover, the statement can describe several aspects: the desirability of the statement (*Appreciatif*), the will (*Boulique*), to what extent the statement is considered moral (*Déontique*), or the degree of uncertainty assigned to the statement or its negation (*Epistémique*).

Orientation captures the hatred-related orientation of speech, with three categories: hateful and not-hateful, and undetermined for ambiguous cases. In addition to those concepts, three relations enrich the structure of the ontology of hate: *hasTarget*, *hasOrientation* and *hasContext* capture the target, the orientation and the context of a statement, respectively. The ontology of hate is represented in OWL [146] and is composed of 61 concepts structured on a 8-levels hierarchy; the model also has 4 *ObjectProperties*.

Hate detection and limitations of the model

The ontology highlights the main concepts for hate detection and fig. 7 shows two examples where hatred message are detected thanks to the identification of hate-specific targets and actions, respectively.




Figure 7: Examples of hate detection

The main purpose of this contribution is to build a formal model for online hate analysis and detection. Practical applications for hate detection will require annotation procedures, through which concepts of the ontology will be related to textual segments. However, practical cases using the ontology of hate should take into account its limitations and the fact that hate speech is contextually embedded, so that comments that in one community are perceived and hateful or offensive are not so in another community [62].

From a practical perspective, detection of specific targets, identification of implicit targets, analysis of irony and sarcasm are the main technical challenges for effective hate speech detection. The following examples, collected from Twitter from February to May 2020, illustrate those challenges.

Political opinions are not part of hate speech characteristics and the example, despite its inherent violence, should thus not be considered as a hateful speech: *Oui! MLP.Philipe Vardon et GÉNÉRATION IDENTITAIRE. Contre la pourriture Estrosi UMPS - LR..*

The presence of hate target but without violent action should be considered, again, as non hateful. *Quand il s'agit d'entretenir les étrangers parasites on est toujours sûr de trouver la gauche.*

The difficulty of detecting implicit or tacit targets is illustrated by the next example: *Attendrons-nous qu'il soit trop tard ? Dans ma profession de foi écrite pour la même campagne législative de 1975 : Les Français ne sont ni xénophobes ni racistes. [...] Mais ils ne supportent plus que la France soit colonisée, exploitée, terrorisée.*

Moreover, sarcasm and irony potentially flip the hateful/no hateful orientation of sentences, as shown fereatfer: *Que se passe-t-il ? Allah est grand ou il n'est pas grand ? Faudrait savoir tout de même ! Il a quand même créé la terre plate !*

Examples above underline the importance of context when detecting hateful and non-hateful contents. Practical applications rely often on the joint utilisation of ontologies and classification or learning techniques to achieve better results. In addition, conceptual models for subjectivity and hate content description are inherently incomplete and constrained to specific standpoints. In practice, the model has to be experienced as valid and

fit for purpose. The next contribution illustrates how ontology-based procedures can be combined with classification algorithms for social data analysis.

3 Illustration of ontology-driven social data analysis

3.1 Detection of trends in online data

One of the applications of models described previously was the detection of trends from social data. Trend analysis was carried out by investigating topics and opinions within several collections of tweets. The rationale behind this distinct analysis is that opinions expressed towards topics are good indicators of the support behind ideologies and ideas. Our original contributions stem from the combination of unsupervised classification techniques for topics detection and the use of ontology-driven techniques to detect fined-grained appraisal categories expressing support, deny, rejection or endorsement.

A hybrid architecture for trend analysis

In order to detect trends and to identify the emerging topics associated with those trends we adopt the following approach:




Figure 8: Hybrid architecture for trend analysis

- Several collections of tweets are crawled and classified thanks to unsupervised techniques;
- Topics are detected within each cluster as to pool tweets into more homogeneous groups and show which topics are the most dominant within each cluster;
- The ontology of appraisal categories is then used to detect appraisal categories in tweets;

- A trend-detection strategy is adopted based on appraisal categories to infer categories detected in topics and to aggregate trends at topic level from appraisals identified at tweet level.

Unsupervised learning and feature engineering

The method adopted for data clustering is DBSCAN (Density-Based Spatial Clustering of Applications with Noise) [159]. This technique groups data located in the region with high density of the data space to belong to the same cluster and it discovers clusters with arbitrary shape. The technique also increases cluster homogeneity by filtering out noise.

The DBSCAN algorithm basically requires 2 parameters:

- `eps`: specifies how close points should be to each other to be considered a part of a cluster. If the distance between two points is lower or equal to `eps` value, those points are considered neighbors.
- `minPoints`: the minimum number of points to form a dense region. For example, if is the `minPoints` parameter is set as 5, then at least 5 points are needed to form a dense region.

The main advantage of DBSCAN is that the method is not restricted to a set number of clusters during initialization, although the algorithm will determine a number of clusters based on the density of a region.

For this work, DBSCAN has been applied iteratively on disjoint dataset portions and the original dataset is clustered at the first level. Then, tweets labelled as outliers in the previous level are re-clustered at each subsequent level. In an attempt to discover representative clusters within datasets, the methods avoid clusters containing few tweets. The number of tweets labelled as outliers and thus unclustered is also limited, in order to consider all different posted items. Through addressing these issues, DBSCAN parameters were properly set at each level by using the Elbow heuristic [155].

Several research surrounding Density-based algorithms in Twitter mining show that they are suitable for clustering unstructured data [14]. However, unsupervised classification puts forth a particular challenge because the algorithm groups data into similar categories, without requiring prior understanding of the groups content. In the context of this work clusters can be of different topics and Latent Dirichlet Allocation(LDA) [128] was used to extract the topics from the vocabulary of previously defined clusters.

LDA is a probabilistic model developed to describe collections of discrete data such as textual corpora composed of words. From a technical standpoint, LDA is a three-level hierarchical Bayesian model, in which each item of a collection is modeled as a finite mixture over an underlying set of topics. Iteratively, each topic is also modeled as an infinite mixture over an underlying set of topic probabilities. In the context of text modeling, the topic probabilities association provides an explicit representation of a document.

Although many clustering models restrict a document to being associated with a single topic, under LDA documents can be associated with multiple topics, as the topic node is sampled repeatedly within the document. The method is also unsupervised and topics are detected without needing to define tags or train data beforehand.

Experimental settings

Datasets: Data acquisition was done by crawling several sets of tweets with specific keywords and additional constraints to select only posts written in English. This phase was carried out as a straightforward step and relies only on data content of both keywords and posts and the ability of API used to search the social network. Three distinct data sets were collected based on query terms, without manual processing or tagging:

- Data set 1: *Brexit*
- Data set 2: *Western values, White supremacy*
- Data set 3: *Security, Hate occident*

Preprocessing: This step involves a series of techniques aiming to optimize the performance of experiments. Prior to conducting experiments, tweets are broken down into sentences and then the preprocessing includes: tokenization, lower casing all tokens, removing images, punctuation and stop words. Furthermore usernames, retweets, duplicates and web references are removed as well.

Tab. 4 shows the number of tweets before and after the preprocessing step and also highlights a significant drop in tweets after the preprocessing.

Number of	Data set 1	Data set 2	Data set 3
Original tweets	103 764	51 000	320 000
Preprocessed tweets	12 057	64 203	124 384

Table 4: Data sets

The rationale behind this cleaning step is that tweets often get large number of retweets and keeping them in the dataset will produce large clusters containing redundant tweets, and it will reinforce false patterns.

Numerical representation : The initial word representations of tweets are initialized with a 300 dimension GloVe vectors which are pretrained [101]. Data sets are split into three parts, with 60/20/20 ratios for training, validation and test sets. Learning algorithms were implemented and run in Python.

Semantic annotation for appraisal detection: attaches additional information to tweets based on their content. Semantic annotation is performed automatically, by using lexical similarities measures that associate a real number to a pair of words and offers a measure of the degree to which two words are similar. For this work, lexical similarities are used to assign words to concepts of the appraisal ontology.

Experimental protocol: Each tweet is treated as an individual document for classification and appraisal annotation purpose. After classification, tweets belonging to the same cluster are combined into one document to both extract dominant topics with LDA and aggregate individual appraisal categories into a general trend at cluster level. The reason for pooling tweets into one document is to overcome the sparsity of twitter data for specific topics or appraisal categories.

Results and evaluation

Results of clustering and topics detection : Parameters for unsupervised clustering were tuned by using the Elbow heuristic method. This method takes into account the variance, explained as a function of the number of clusters to be build and identifies graphically an elbow threshold highlighting to the number of clusters corresponding to the optimal value of variance. Mathematically,

$$W_k = \sum_{r=1}^k D_r / n_r \quad (1)$$

Where k is the number of clusters, n_r is the number of points in cluster r and D_r is the sum of distances between all points in a cluster:

$$D_r = \sum_{i=1}^{n_r-1} \sum_{j=1}^{n_r} \|d_i - d_j\|^2 \quad (2)$$

Tab. 5 shows the number of clusters created according to the Elbow criterion and the percentage of clustered data for each data set.

	Data set 1	Data set 2	Data set 3
Number of clusters	2	5	10
Clustered data (%)	97	81	76

Table 5: Number of clusters and percentage of clustered data

In addition, topics coherence scores were used to evaluate the coherence of clusters and the results were visualized by using the pyLDAvis library ¹¹. Intuitively, the coherence measure captures the precision of the clustering, and values under 100 are due to the fact that the categorization of tweets is not crisp, and some tweets belong to several clusters.

For the first data set, the clustering algorithm identified 2 clusters illustrated in fig. 9, and topic detection highlighted several words associated to each cluster.




Figure 9: Clusters of data set 1 (3D)

¹¹<https://pypi.org/project/pyLDAvis/>

Among the topics, a main topic clearly highlights *Theresa May* as dominant named entity. The number of clusters for the second data set is 5. Tab. 6 shows a topic as build by LDA. Topics detection clearly identifies topics related to white supremacy, although the overall data set was collected with 2 keywords.

topic 1	supremacy	commoner	sacrifice	opinion
topic 2	protect	surprise	person	promote
topic 3	patriarchy	racism	child	never

Table 6: Example of topics and most frequent words

The processing of the last set provides 10 clusters, see fig. 10.




Figure 10: Clusters for data set 3 (2D)

This data set is the most heterogeneous, due to the keyword *security* that gathers a large variety of tweets. The distribution of topics detected within is also sparse.

Appraisal annotation : After the clustering phase, the appraisal ontology was used to detect instances of appraisal categories at tweet level. Thus, tweets were annotated through the ontology of appraisals. Tab. 7 shows the analysis of data sets in terms of numerical distribution of subjective and objective tweets.

	Data set 1	Data set 2	Data set 3
Annotated data (%)	96	83	79

Table 7: Percentage of annotations

Numerical values show a high percentage of subjective tweets for all datasets analyzed, regardless of the keyword used to collect. The trend is still the same at cluster level, with clusters having similar values for the ration of subjective/ objective content.

Fig. 11 illustrates the distribution of main appraisal categories for the three data sets.

The main appraisal categories detected are: *Attitude* (as a general concept), *affect* (specific concept under attitude), and *support*, *deny*, *disapproval* and *submission*, all specific concepts under *Engagement* concept.

Thus, data set1 collected with Brexit keyword, appears related to *affect*, *deny*, *disapproval* and *submission* concepts while being less correlated with concepts under *Attitude* and *Support*. The same associations appear when analyzing individually each cluster.




Figure 11: Distribution of main appraisal categories in data sets

Data set 2 is strongly correlated with *Deny* concept, and has an almost equivalent distribution of *affect*, *attitude*, *support*, *submission* and *disapproval* categories.

Data set 3 has a strong affect component, and is strongly correlated with *Attitude* and *Affect*. Under *Engagement* system, *Disapproval* is the most relevant concept associated to this data set. Fig. 12 shows the distribution of positive and negative polarities, with *high* and *low* force.




Figure 12: Distribution of polarities in data sets

Regarding the orientation of tweets, there is a good representation of negative-oriented tweets that account for more than 50% for each collection. Although the negative orientation is prevalent, the majority of tweets are of moderate impact, and tweets having low or high impact are underrepresented for all data sets.

Evaluation of trend detection Quantitative analysis by manually inspecting the quality of output was also undertaken to validate the results provided by the overall approach. Ground truth was created for the data set 1 by analyzing the type and polarity of annotation in the light of Google news headlines, see tab. 8.

True positives (Tp)	7211
Trues negatives (Tn)	972
False positives (Fp)	913
False negatives (Fn)	2300
Undetermined (U)	661

Table 8: Manual validation and ground truth in %

In the light of those values, the quality of trend detection is estimated with *Precision*,

Recall and Accuracy.

Values of metrics (%) are: Precision = 88, Recall = 75 and Accuracy =72.

Statistical values discussed above give a general overview of the association between data sets, appraisal and their orientations as build thanks to the clustering algorithms and semantic annotation. In addition, a manual analysis of results sheds light on a number of findings that can be useful in better understanding such diverse associations of subjective states and real-life topics.

Thus, through the appraisal pooling strategy, positive polarities are associated with Theresa May, whilst negative categories often included references to political system, recession and war. Similarly, Brexit was frequently mentioned in tweets associated later to topics that belonged to positive polarities, including terms such opinion, person and symbol.

From a different perspective, topics such as *supremacy* were associated with negative emotions such as deny and disapproval.

And finally, within the last data set, terms such as protect, surprise and patriarchy were closely associated to categories such as submission, deny and affect.

4 Related work

Social data analysis has been actively studied in the past few years, and more specifically Twitter became an increasingly popular source of data, with mainly machine and deep learning models being developed for this task [41]. The studies keep the distinction between factual and subjective aspects and focus either on topics [92], [66] and narrative detection [26], [64] or subjectivity analysis [89] including the detection of sentiments [73] or opinions [90]. Currently, most work in the area of social data processing focuses on subjectivity analysis, regardless of the entities, targets or topics mentioned in the content. Most of the algorithms detect sentiments or opinion by assuming a known target. This disjoint analysis is a limitation of those approaches as meaningful clues hidden in online data are often a combination of topics and subjective aspects and their identification involves analysis of emotions conveyed towards specific topics. Only a very few studies have considered the joint analysis of topics and subjectivity. Among them, Schoene and de Mel [1] investigate the correlation of topics and emotions while Vijayaraghavan and colleagues addressed the classification of topics along with the sentiment [80]. Following a similar research line, Schmitt and colleagues employed an end-to-end trainable neural network to detect aspects and classify polarities jointly.

From a different perspective, three main research directions have been developed to tackle online hate : building databases and resources, detecting hate versus non hate content and classifying specific subcategories of hate messages.

Some studies aimed at building repositories of hatred messages to feed technological blocks. Among the resources available online, Hurltlex ¹² is a multilingual computational lexicon of hate words and Hatebase ¹³ is an online repository of structured, multilingual and usage-oriented hate speech. The base is structured thanks to a vocabulary composed

¹²<https://vocabularyserver.com/hurtlex/>

¹³<https://www.hatebase.org/>

of eight categories, namely : archaic, class, disability, ethnicity, gender, nationality, religion, and sexual orientation. Another collection available online is the Kaggle's Toxic Comment Classification Challenge dataset ¹⁴, consisting of around 150k Wikipedia comments annotated for toxic behaviour. More specifically, a date set of hateful comments from a white supremacy forum is described in [37]. Hatebase and its vocabulary was used by Serra and colleagues to build classifiers for hate speech [63]. A variety of techniques were used for this binary classification of hateful non hateful content, including machine learning [75], classification with a mix of features such as images and emojis [21]. Several approaches also tackled the classification of hate content according to finer categories, such as abusive [12] or aggressive [40], language, racism [62] and extremist contents[48].

There are a number of recent works on the definition of language-specific methods for social data analysis, but one of the main problem researchers face when dealing with language aside from English is the lack of corpora, lexical resources and ontologies. This holds also for French and translating English resources does not offer good results. Ghorbel and Jacot [123] translated English SentiWordNet entries into French and showed that even if the translation of entries is correct, some parallel words do not always share the same polarities or meaning across both languages, simply due to differences in common usage of words.

Our studies on social data analysis provide a comprehensive description of affect states in online data. In particular, two key aspects are taken into account: (i) considering a cognitive background that describes features of subjectivity and (ii) modeling this framework and highlighting the linguistic dimension of the concepts.

5 Lessons learnt

In this chapter I addressed some aspects of social data analysis, namely the detection of subjective states such as general opinions or hateful comments. The cyberspace is a valuable information source for defence and security applications and many solutions have been developed. These approaches achieve reasonable accuracy on limited tasks, such as sentiment analysis, but encounter difficulties in processing a mix of multi-domain data and the use of manually tagged inputs have been plaguing their robustness [57].

The analysis of several use-cases using social data in conjunction with real life events (terrorist attacks, crisis) showed the limitations of such applications to detect threats in an effective manner [3]. Reasons are manifold: social data are vast, noisy, unstructured, inherently dynamic and heterogeneous in nature and processing this type of data is challenging for traditional data mining approaches that are often too slow and expensive, rely on sample sizes, and come with biases leading to errors [108]. There is currently a crucial need to carefully consider the evaluation and scalability aspects. Evaluation protocols are still considering manual evaluation on sample and sampled of data are needed to be built for crossed evaluation.

In addition to evaluation protocols, achieving better results by designing hybrid approaches for social data analysis is a research line worth to be investigated. The first

¹⁴<https://www.kaggle.com/c/jigsaw-toxic-comment-classification-challenge/data>

hybridization line can combine social and technical inputs: indicators collected from experts who have experience with solving a rather narrow and rare empirical problem and help designing automatic procedures for those rare cases.

Another hybridization line can combine learning and semantic-driven approaches, such as implementing flexible learning approaches able to take into account inputs inferred from knowledge or rules and ensuring this way the dynamic and domain-adaptive exploration of contents.

The overall conclusion of the analysis conducted in this chapter is that social stream exploration is currently a hard problem in several key ways and addressing these aspects can significantly improve the accuracy of techniques.

6 Collaborations and dissemination of results

This chapter presents my contributions to the problematic of social data analysis. These contributions have been done in the context of several projects and collaborations.

PRACTICIES - H2020, Security

The analysis of extremist content online started within the frame of PRACTICIES, an European (H2020) project led by Prof. Seraphin Alava, from University Toulouse 2, Jean Jaures. For this project, I contributed to the analysis of online propaganda and I had collaborations with **S raphin Alava**, on social and educational aspects, and **Claudine Moise**, from University Grenoble Alpes, on sociolinguistics aspects [35].

FLYER - French Research Agency and the French MoD

The development of dynamic techniques for effective online investigation was the goal of FLYER, a research project funded by the French Research Agency. I was the scientific coordinator of the project and my close collaboration with **Delphine Battistelli**, from MODYCO, University of Paris - Nanterre, allowed the development of linguistically-grounded methods and models [32], [2]. The implementation of techniques was done thanks to collaborations with **Fr d rique Segond** from INRIA and **Romain Campigotto** from BERTIN IT.

On aspects related to ethics, which were particularly important for this research topic, I collaborated with my colleague **Catherine Tessier**, from ONERA Toulouse.

ROSARIO

ROSARIO was a research project funded by ONERA and coordinated by **Laurence Cholvy**. ROSARIO gave me the opportunity to collaborate with colleagues from ONERA's Toulouse research centre: **Claire Saurel**, **Olivier Poitou** and **Martin Adelantado** [47].

LEMON – coordination action, MADICS, CNRS

LEMON is a coordination action funded by MADICS, a cluster of the French National Research Centre. Together with Delphine Battistelli we coordinate LEMON, for which we have organized several workshops on social data analysis.

ETUR use case on social media analysis

As part of my participation to the ETUR working group, I coordinated together with Claire Laudy from Thales TRT, France an use case dedicated to social media analysis [4].

NATO Research track Exploitation of Cyberspace for Intelligence (IST-159 RTG)

I was appointed to integrate this NATO research task in January 2018. This task offers a place to exchange with researchers from DRDC Canada, FKIE, Germany NATO StratCom (Latvia), the Baltic Institute of Technology (Lithuania) and the university of Helsinki (Finland).

Supervision of students

Several master students contributed to this research topics: Emmanuelle Kellodjoue modeled the appraisal ontology in English. This model was then adapted to french content analysis by Adrien Legros. Aurélien Mascaro implemented the approach for trend detection and Océane Fourquet used jointly the appraisal ontology and knowledge graphs to extract semantic information from social data. Cyril Bruneau contributed to the ontological analysis of online hate.

Dissemination of results

It should be noted that the dissemination of results presented in this chapter was affected by restrictions imposed by the European Commission, given the sensitivity of the topics addressed, and two reports [36] and [35] cannot be released to the scientific community. In addition to scientific articles, the dissemination also included :

- An invited tutorial entitled *Social data analysis for intelligence* presented at the International Conference on Cognitive and Computational Aspects of Situation Management, CogSIMA 2019;
- An invited talk in the context of the *Defence Security Cyber* summer school, Bordeaux 2018;
- An invited talk to the Center for Cyberspace Analysis, an operational service of the French Intelligence Agency, Creil, 2019;
- A paper and presentation for the NATO workshop *Big Data Challenges: Situation Awareness and Decision Support* , Budapest 2019;

- Two ALIAS workshops organized in conjunction with major French conferences in 2018 and 2019.

Main publications of the chapter

- [2] Delphine Battistelli, Cyril Bruneau, and Valentina Dragos. “Building a formal model for hate detection in French corpora”. In: *24th International Conference in Knowledge Based and Intelligent Information and Engineering Systems, KES 2020*. Procedia Computer Science. Elsevier, 2020.
- [3] Valentina Dragos, Bruce Forrester, and Kellyn Rein. “Is hybrid AI suited for hybrid threats? Insights from social data analysis”. In: *Proceedings of the 23rd International Conference on Information Fusion*. IEEE. 2020.
- [32] Valentina Dragos, Delphine Battistelli, and Emmanuelle Kelodjoue. “Beyond sentiments and opinions: exploring social media with appraisal categories”. In: *2018 21st International Conference on Information Fusion (FUSION)*. IEEE. 2018, pp. 1851–1858.

Chapter 5

Conclusion and perspectives

This project presents an overview of my contributions to the domain of knowledge and semantic technologies for situation awareness. Those contributions have been obtained by combining standpoints from computer science, linguistics and sociology and have systematically involved close collaborations with domain experts and operators in the field. Approaches developed aimed at designing knowledge models for intelligence analysis and maritime surveillance, assessing the quality of reported information and defining semantic-based procedures for cyberspaces exploration.

Challenges in each of these domains are still numerous. The huge amount of complex data produced on social platforms open the door to several open research questions directly related to the big data paradigm: how to store, index, query and efficiently analyze the huge and highly distributed amounts of data concretely produced on social platforms and how to make sense out of them in order to better support decision making?

Several complexity questions are key points to be addressed in the future, considering content characterization and understanding salient data features to design efficient solution for situation awareness. Advances in building accurate situations depend – and may have impacts on – progress made in other communities such as natural language processing, socio-linguistics, social networks, supervised and unsupervised learning or more generally algorithmic and graphs to manage complex data.

There are several research directions that I am interested to investigate in the next three to five years.

Finer characterization of online content

This perspective is in line with the approaches developed for ROSARIO and PRACTICES projects. Contributions on this topic will allow overcoming the limitations describing content with keywords, in order to provide in-depth understanding of online content by highlighting characteristics of propaganda, violence or attitudes.

A first contribution on this research line is planned within the framework of the FLYER project which aims to detect extremist content on the Internet. Collaborations are also taking place within the framework of the NATO research group *Social Media Exploitation for Operations in the Information Environment*.

Hybrid techniques for artificial intelligence

Developing hybrid methods for artificial intelligence is another research line tackled in relation with the RIAD and FLYER projects. The intention is to develop new processing paradigm able to go beyond the current limitations of supervised learning approaches. The intuition is that a static set of features is not appropriate to analyze dynamic content, and potentially can fail in detecting new topics. It becomes then necessary to update all the set of features by taking into account two sources: formal knowledge, capturing the evolution of the domain according to experts and attributes specific to the data, in the form of relevant associations of words extracted from data. The hybridization of learning and semantic-based approaches can offer a solution to tackle the analysis of highly dynamic contents.

Analysis of uncertainty for situation understanding

In the maritime domain, using the criteria defined by the ETUR working group can improve the description of maritime situations by indicating the sources and types of uncertainty. Investigating the propagation of uncertainties is also a direction to be considered since maritime situations are built by merging various information types at different moments in time.

For the aeronautics domain, the analysis of uncertainties will be associated with the integration of several knowledge bases, knowledge models and schemata that were built independently in order to describe different aspects of the field, such as entities, technical problems and failures.

Ethical aspects of developing systems for defence and security applications

Research actions on information processing do not consider whether and with what capabilities defence and security systems should be developed, but those actions have to assess the ethical dimension of how those systems are designed, implemented and used. This task systematically includes investigating technical options and limitations of those systems, researching methods for their use and control, advising users, evaluate the risk associated to their development and estimate future threats. By considering ethical aspects, situation awareness will offer not only a basis for decision making, but also a guarantee that the support is provided in an ethically and societally responsible fashion.

Résumé

Ce projet de recherche présente des méthodes et outils que j'ai développés au cours des dix dernières années ainsi que les perspectives de mes activités scientifiques. L'objectif de ces activités était d'intégrer des connaissances dans les chaînes de traitement de l'information pour les applications de sécurité et défense. Mes contributions ont abordé plusieurs aspects. Premièrement, deux approches ont été développées pour évaluer la qualité des informations fournies par les humains. Ces solutions utilisent des ontologies pour détecter les incohérences et les contradictions, ou exploitent les propriétés des sources d'information pour attribuer un degré de confiance aux items. Deuxièmement, nous avons étudié l'utilisation de la sémantique pour la recherche d'informations, la réutilisation des connaissances expertes et la fusion des données hétérogènes. Troisièmement, nous avons développé des modèles et algorithmes pour l'exploitation des sources ouvertes et des données sociales dans le contexte des applications de défense et de sécurité. Ces approches ont été développées en collaboration avec des collègues du monde académique et industriel et incluent plusieurs contributions de stages de master et postdoctoraux. Les travaux ont été menés dans le cadre de deux projets de recherche que j'ai coordonnés et de plusieurs projets de recherche financés par l'ANR, DGA, CNRS et la Commission européenne. Les projets impliquaient des partenaires académiques et industriels, tels que George Mason University, (USA), DRDC (Canada), FKIE (Allemagne), NATO CMRE (Italie), Université Paris Nanterre, Université Toulouse 2 (France), Thales TRT et Airbus DS (France)).

Mots-clés: Cadres sémantiques, ontologies, intelligence artificielle, fusion de haut-niveau, défense et sécurité

Abstract

This research project outlines some methods and tools that I have developed during the last ten years and the perspectives of my scientific activities for the near future. The goal of those activities was to integrate knowledge into various processing chains implemented for defence and security applications. My contributions focused on several aspects. First, assessing the quality of information provided by humans, for which two main approaches were developed relying on either the semantics of ontologies to detect inconsistency and contradictions or on the properties of sources in the reporting chain to assign a degree of trust for every information item. Second, we have investigated the use of semantics as a backbone for information retrieval, reuse of experiences and heterogeneous data fusion. We developed approaches that contribute to the overcoming problem of dynamically integrating resources, user feed-back, sensor data and observations, which is a critical aspect of situation awareness today. Third, we addressed the exploitation of open sources and social data in the context of defence and security applications. Those approaches have been developed in collaboration with several fellow researchers and include contributions from several master thesis and post docs. The work was carried out in the framework of two research projects that I coordinated and several research projects to which I contributed, with funding from the European Commission, ANR and CNRS. The projects involved academic and industrial partners, such as George Mason University (USA), DRDC (Canada), FKIE (Germany), NATO CMRE (Italy), University Paris Nanterre and University Toulouse 2 (France), Thales and Airbus (France).

Keywords: Semantic frameworks, ontology, high-level information fusion, situation awareness, defence and security

References

- [1] Annika Schoene and Geeth de Mel. “Pooling Tweets by Fine-Grained Emotions to Uncover Topic Trends in Social Media”. In: ().
- [2] Delphine Battistelli, Cyril Bruneau, and Valentina Dragos. “Building a formal model for hate detection in French corpora”. In: *24th International Conference in Knowledge Based and Intelligent Information and Engineering Systems, KES 2020*. Procedia Computer Science. Elsevier, 2020.
- [3] Valentina Dragos, Bruce Forrester, and Kellyn Rein. “Is hybrid AI suited for hybrid threats? Insights from social data analysis”. In: *Proceedings of the 23rd International Conference on Information Fusion*. IEEE. 2020.
- [4] Claire Laudy and Valentina Dragos. “Use cases for social data analysis with URREF criteria”. In: *Proceedings of the 23rd International Conference on Information Fusion*. IEEE. 2020.
- [5] Sanju Mishra and Sarika Jain. “Ontologies as a semantic model in IoT”. In: *International Journal of Computers and Applications* 42.3 (2020), pp. 233–243.
- [6] Hui Yie Teh et al. “Sensor data quality: a systematic review”. In: *Journal of Big Data* 7.1 (2020), p. 11.
- [7] Mashaël Almoqbel and Songhua Xu. “Computational Mining of Social Media to Curb Terrorism”. In: *ACM Computing Surveys (CSUR)* 52.5 (2019), pp. 1–25.
- [8] Ernst-Josef Behmer et al. “Ontology Population Framework of MAGNETO for Instantiating Heterogeneous Forensic Data Modalities”. In: *IFIP International Conference on Artificial Intelligence Applications and Innovations*. Springer. 2019, pp. 520–531.
- [9] Erik P Blasch et al. “Uncertainty Ontology for Veracity and Relevance”. In: *2019 22th International Conference on Information Fusion (FUSION)*. IEEE. 2019, pp. 1–8.
- [10] Valentina Dragos, Jean Dezert, and Kellyn Rein. “Assessment of Trust in Opportunistic Reporting using Belief Functions”. In: *22nd International Conference on Information Fusion, Ottawa, Canada*. IEEE. 2019, pp. 208–216.
- [11] Valentina Dragos et al. “Entropy-Based Metrics for URREF Criteria to Assess Uncertainty in Bayesian Networks for Cyber Threat Detection”. In: *22nd International Conference on Information Fusion, Ottawa, Canada*. IEEE. 2019, pp. 1–8.

- [12] Antigoni Maria Founta et al. “A unified deep learning architecture for abuse detection”. In: *Proceedings of the 10th ACM Conference on Web Science*. 2019, pp. 105–114.
- [13] Armin Haller et al. “The modular SSN ontology: A joint W3C and OGC standard specifying the semantics of sensors, observations, sampling, and actuation”. In: *Semantic Web 10.1* (2019), pp. 9–32.
- [14] Ryosuke Harakawa et al. “Consensus Clustering of Tweet Networks via Semantic and Sentiment Similarity Estimation”. In: *IEEE Access* 7 (2019), pp. 116207–116217.
- [15] Andrew Iliadis. “The Tower of Babel problem: making data make sense with Basic Formal Ontology”. In: *Online Information Review* (2019).
- [16] Daniel Irwin and David Mandel. “Standards for Evaluating Source Reliability and Information Credibility in Intelligence Production”. In: *Irwin, D. & Mandel, DR (accepted conditionally). Standards for evaluating source reliability and information credibility in intelligence production. In DR Mandel (Ed.), Assessment and Communication of Uncertainty in Intelligence to Support Decision Making: Final Report of Research Task Group S* (2019).
- [17] Krzysztof Janowicz et al. “SOSA: A lightweight ontology for sensors, observations, samples, and actuators”. In: *Journal of Web Semantics* 56 (2019), pp. 1–10.
- [18] Myung-Joong Jeon et al. “Relationship between UAVs and Ambient objects with threat situational awareness through grid map-based ontology reasoning”. In: *International Journal of Computers and Applications* (2019), pp. 1–16.
- [19] Michael J Mazarr et al. *Hostile Social Manipulation Present Realities and Emerging Trends*. Tech. rep. RAND Corporation Santa Monica United States, 2019.
- [20] Sanju Mishra and Sarika Jain. “Towards a Semantic Knowledge Treasure for Military Intelligence”. In: *Emerging Technologies in Data Mining and Information Security*. Springer, 2019, pp. 835–845.
- [21] Juan Carlos Pereira-Kohatsu et al. “Detecting and Monitoring Hate Speech in Twitter”. In: *Sensors* 19.21 (2019), p. 4654.
- [22] Kellyn Rein. “Fact, Conjecture, Hearsay and Lies: Issues of Uncertainty in Natural Language Communications”. In: *Information Quality in Information Fusion and Decision Making*. Springer, 2019, pp. 155–179.
- [23] Galina L Rogova. “Information Quality in Fusion-Driven Human-Machine Environments”. In: *Information Quality in Information Fusion and Decision Making*. Springer, 2019, pp. 3–29.
- [24] Galina L Rogova and Lauro Snidaro. “Quality, Context, and Information Fusion”. In: *Information Quality in Information Fusion and Decision Making*. Springer, 2019, pp. 219–242.
- [25] K Srinivasa and P Santhi Thilagam. “Crime base: Towards building a knowledge base for crime entities and their relationships from online news papers”. In: *Information Processing & Management* 56.6 (2019), p. 102059.

- [26] Jichuan Zeng et al. “What you say and how you say it: Joint modeling of topics and discourse in microblog conversations”. In: *Transactions of the Association for Computational Linguistics* 7 (2019), pp. 267–281.
- [27] Abdulwhab Alkharashi and Karen Renaud. “Privacy in crowdsourcing: a systematic review”. In: *International Conference on Information Security*. Springer. 2018, pp. 387–400.
- [28] Miguel Alberto Becerra, Karla C Alvarez-Uribe, and Diego Hernán Peluffo-Ordoñez. “Low data fusion framework oriented to information quality for BCI systems”. In: *International Conference on Bioinformatics and Biomedical Engineering*. Springer. 2018, pp. 289–300.
- [29] Adrian N Bishop et al. “Spatio-temporal tracking from natural language statements using outer probability theory”. In: *Information Sciences* 463 (2018), pp. 56–74.
- [30] Sarah B Bouchat. “InDetermination: Measuring Uncertainty in Social Science Texts”. In: (2018).
- [31] Paulo C. G. Costa et al. “URREF: Uncertainty representation and reasoning evaluation framework for information fusion”. In: *Journal of Advances in Information Fusion* 13 (2018), pp. 137–158.
- [32] Valentina Dragos, Delphine Battistelli, and Emmanuelle Kelodjoue. “Beyond sentiments and opinions: exploring social media with appraisal categories”. In: *2018 21st International Conference on Information Fusion (FUSION)*. IEEE. 2018, pp. 1851–1858.
- [33] Valentina Dragos, Jean Dezert, and Kellyn Rein. “Tracking uncertainty propagation from model to formalization: Illustration on trust assessment”. In: *Journal of Advances in Information Fusion* 13 (2018), pp. 198–216.
- [34] Valentina Dragos and Sylvain Gatepaille. *Semantic mediation for heterogeneous fusion and enhanced situation assessment*. Workshop on AI and Cognitive Technologies for Radar, Comms and EW, Workshop organized by the European Defense Agency, December, 2018. Dec. 2018.
- [35] Valentina Dragos, Claudine Moise, and Lluís Botifoll. *Opinion mining from open sources*. RT 1/26335 - DTIS - PALAISEAU. 2018.
- [36] Valentina Dragos, Claudine Moise, and Lluís Botifoll. *Semantic approaches to analyse radicalized content: challenges and novel solutions*. RT 2/26335 - DTIS - Palaiseau, 2018.
- [37] Ona de Gibert et al. “Hate speech dataset from a white supremacy forum”. In: *arXiv preprint arXiv:1809.04444* (2018).
- [38] Andrew Janes. “An exploratory study In challenges faced by intelligence leaders and practitioners from goal and structural conflicts”. In: (2018).
- [39] Alexander Kott. “Challenges and characteristics of intelligent autonomy for internet of battle things in highly adversarial environments”. In: *2018 AAAI Spring Symposium Series*. 2018.

- [40] Ritesh Kumar et al. “Benchmarking aggression identification in social media”. In: *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*. 2018, pp. 1–11.
- [41] Saif Mohammad et al. “Semeval-2018 task 1: Affect in tweets”. In: *Proceedings of the 12th international workshop on semantic evaluation*. 2018, pp. 1–17.
- [42] Manuela Sanguinetti et al. “An italian twitter corpus of hate speech against immigrants”. In: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. 2018.
- [43] Feifei Shi et al. “A survey of data semantization in internet of things”. In: *Sensors* 18.1 (2018), p. 313.
- [44] Leslie F Sikos et al. “Automated reasoning over provenance-aware communication network knowledge in support of cyber-situational awareness”. In: *International Conference on Knowledge Science, Engineering and Management*. Springer. 2018, pp. 132–143.
- [45] Patrick Wang, Rafael Angarita, and Ilaria Renna. “Is this the era of misinformation yet: combining social bots and fake news to deceive the masses”. In: *Companion Proceedings of the The Web Conference 2018*. International World Wide Web Conferences Steering Committee. 2018, pp. 1557–1561.
- [46] Amine Abdaoui et al. “Feel: a french expanded emotion lexicon”. In: *Language Resources and Evaluation* 51.3 (2017), pp. 833–855.
- [47] Martin Adelantado et al. *Contributions to Open sources analysis and Opinion dynamics modeling and simulation*. Séminaire TIM (Taitement de l’Information Multimodale) organisé par la DGA, Paris. June 2017.
- [48] Pinkesh Badjatiya et al. “Deep learning for hate speech detection in tweets”. In: *Proceedings of the 26th International Conference on World Wide Web Companion*. 2017, pp. 759–760.
- [49] Arne Bröring et al. “Enabling IoT ecosystems through platform interoperability”. In: *IEEE software* 34.1 (2017), pp. 54–61.
- [50] Valentina Dragos. “Detection of contradictions by relation matching and uncertainty assessment”. In: *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 21st International Conference KES-2017*. Vol. 112. Procedia Computer Science. Elsevier, 2017, pp. 71–80.
- [51] Valentina Dragos and Sylvain Gatepaille. “On-the-fly integration of soft and sensor data for enhanced situation assessment”. In: *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 21st International Conference KES-2017*. Vol. 112. Procedia Computer Science. Elsevier, 2017, pp. 1263–1272.
- [52] Mica R Endsley. “Toward a theory of situation awareness in dynamic systems”. In: *Situational Awareness*. Routledge, 2017, pp. 9–42.
- [53] Emilio Ferrara. “Contagion dynamics of extremist propaganda in social networks”. In: *Information Sciences* 418 (2017), pp. 1–12.

- [54] Raffaele Gravina et al. “Multi-sensor fusion in body sensor networks: State-of-the-art and research challenges”. In: *Information Fusion* 35 (2017), pp. 68–80.
- [55] David L Hall and James Llinas. “Multisensor data fusion”. In: *Handbook of multi-sensor data fusion*. CRC press, 2017, pp. 21–34.
- [56] Pierre-Antoine Jean et al. “Gérer l’incertitude lors de l’extraction de relations et lors de l’inférence de nouvelles connaissances”. In: 2017.
- [57] Zheng Li et al. “End-to-End Adversarial Memory Network for Cross-domain Sentiment Classification.” In: *IJCAI*. 2017, pp. 2237–2243.
- [58] Martin Liggins II, David Hall, and James Llinas. *Handbook of multisensor data fusion: theory and practice*. CRC press, 2017.
- [59] Mainack Mondal, Leandro Araújo Silva, and Fabricio Benevenuto. “A measurement study of hate speech in social media”. In: *Proceedings of the 28th ACM Conference on Hypertext and Social Media*. 2017, pp. 85–94.
- [60] Fang-Ping Pai, Lee-Jang Yang, and Yeh-Ching Chung. “Multi-layer ontology based information fusion for situation awareness”. In: *Applied Intelligence* 46.2 (2017), pp. 285–307.
- [61] Lucas Paletta et al. “Towards real-time probabilistic evaluation of situation awareness from human gaze in human-robot interaction”. In: *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM. 2017, pp. 247–248.
- [62] Haji Mohammad Saleem et al. “A web of hate: Tackling hateful speech in online social spaces”. In: *arXiv preprint arXiv:1709.10159* (2017).
- [63] Joan Serra et al. “Class-based prediction errors to detect hate speech with out-of-vocabulary words”. In: *Proceedings of the First Workshop on Abusive Language Online*. 2017, pp. 36–40.
- [64] PK Srijith et al. “Sub-story detection in Twitter with hierarchical Dirichlet processes”. In: *Information Processing & Management* 53.4 (2017), pp. 989–1003.
- [65] Richard M Taylor. “Situational awareness rating technique (SART): The development of a tool for aircrew systems design”. In: *Situational awareness*. Routledge, 2017, pp. 111–128.
- [66] Bo Wang et al. “A hierarchical topic modelling approach for tweet clustering”. In: *International Conference on Social Informatics*. Springer. 2017, pp. 378–390.
- [67] Feng Wang et al. “A semantics-based approach to multi-source heterogeneous information fusion in the internet of things”. In: *Soft Computing* 21.8 (2017), pp. 2005–2013.
- [68] Judson Bandeira et al. “FOCA: A methodology for ontology evaluation”. In: *arXiv preprint arXiv:1612.03353* (2016).
- [69] Valentina Dragos. “Comparative Studies on Using Semantic Filtering for Open Relation Identification.” In: *Meeting Security Challenges Through Data Analytics and Decision Support*. IOS Press, 2016, pp. 69–77.

- [70] Valentina Dragos. “From finding to explaining: Information retrieval to support maritime anomaly analysis”. In: *International Journal of Knowledge and Systems Science* 7 (2016), pp. 40–55.
- [71] Valentina Dragos, Sylvain Gatepaille, and Xavier Lerouvreur. “Refining relation identification by combining soft and sensor data”. In: *2016 19th International Conference on Information Fusion (FUSION)*. IEEE. 2016, pp. 2139–2146.
- [72] Valentina Dragos and Kellyn Rein. “What’s in a message? Exploring dimensions of trust in reported information”. In: *2016 19th International Conference on Information Fusion (FUSION)*. IEEE. 2016, pp. 2125–2132.
- [73] Anastasia Giachanou and Fabio Crestani. “Like it or not: A survey of twitter sentiment analysis methods”. In: *ACM Computing Surveys (CSUR)* 49.2 (2016), p. 28.
- [74] Jukka Jouhki et al. “Facebook’s emotional contagion experiment as a challenge to research ethics”. In: *Media and Communication* 4 (2016).
- [75] Chikashi Nobata et al. “Abusive language detection in online user content”. In: *Proceedings of the 25th international conference on world wide web*. 2016, pp. 145–153.
- [76] Kellyn Rein. “I think it is likely that it might be so... Exploiting Lexical Clues for the Automatic Generation of evidentiality Weights for Information Extracted from English Text”. Dissertation. University of Bonn, 2016.
- [77] Galina L Rogova. “Information quality in information fusion and decision making with applications to crisis management”. In: *Fusion Methodologies in Crisis Management*. Springer, 2016, pp. 65–86.
- [78] Ronald Schroeter and Fabius Steinberger. “Pokémon DRIVE: towards increased situational awareness in semi-automated driving”. In: *Proceedings of the 28th Australian Conference on Computer-Human Interaction*. ACM. 2016, pp. 25–29.
- [79] Elisa Shahbazian and Galina Rogova. *Meeting Security Challenges Through Data Analytics and Decision Support*. Vol. 47. IOS Press, 2016.
- [80] Prashanth Vijayaraghavan, Soroush Vosoughi, and Deb Roy. “Automatic detection and categorization of election-related tweets”. In: *Tenth International AAAI Conference on Web and Social Media*. 2016.
- [81] Min-Jung Yoo et al. “Semantic model for iot-enabled electric vehicle services: Puzzling with ontologies”. In: *2016 IEEE 4th International Conference on Future Internet of Things and Cloud (FiCloud)*. Ieee. 2016, pp. 387–392.
- [82] JM Berger and Jonathon Morgan. “Defining and describing the population of ISIS supporters on Twitter”. In: URL <http://www.brookings.edu/research/papers/2015/03/isis-twitter-census-berger-morgan> (2015).
- [83] Erik Blasch. “One decade of the data fusion information group (dfig) model”. In: *Next-Generation Analyst III*. Vol. 9499. International Society for Optics and Photonics. 2015, p. 94990L.

- [84] Valentina Dragos, Xavier Lerouvreux, and Sylvain Gatepaille. “A critical assessment of two methods for heterogeneous information fusion”. In: *Information Fusion (Fusion), 2015 18th International Conference on*. IEEE. 2015, pp. 42–49.
- [85] Rino Falcone, Alessandro Sapienza, and Cristiano Castelfranchi. “The relevance of categories for trusting information sources”. In: *ACM Transactions on Internet Technology (TOIT)* 15.4 (2015), p. 13.
- [86] Sylvain Gatepaille, Valentina Dragos, and Frédéric Dambreville. “FUTHANE: An Open Framework to Support Heterogeneous Information Fusion”. In: *2015 NATO IST-SET-126 Symposium on "Information Fusion (Hard and Soft) for ISR*. IEEE, Norfolk Virginia, United States. 2015, pp. 21–30.
- [87] Jones Granatyr et al. “Trust and Reputation Models for Multiagent Systems”. In: *ACM Computing Surveys (CSUR)* 48.2 (2015), p. 27.
- [88] Ankur Joshi et al. “Likert scale: Explored and explained”. In: *Current Journal of Applied Science and Technology* (2015), pp. 396–403.
- [89] Ana Carolina ES Lima, Leandro Nunes de Castro, and Juan M Corchado. “A polarity analysis framework for Twitter messages”. In: *Applied Mathematics and Computation* 270 (2015), pp. 756–767.
- [90] Kumar Ravi and Vadlamani Ravi. “A survey on opinion mining and sentiment analysis: tasks, approaches and applications”. In: *Knowledge-Based Systems* 89 (2015), pp. 14–46.
- [91] Lauro Snidaro, Ingrid Visentini, and Karna Bryan. “Fusing uncertain knowledge and evidence for maritime situational awareness via Markov Logic Networks”. In: *Information Fusion* 21 (2015), pp. 159–172.
- [92] Carlos Vicient and Antonio Moreno. “Unsupervised topic discovery in micro-blogging networks”. In: *Expert Systems with Applications* 42.17-18 (2015), pp. 6472–6485.
- [93] Dong Wang, Tarek Abdelzaher, and Lance Kaplan. *Social sensing: building reliable systems on unreliable data*. Morgan Kaufmann, 2015.
- [94] Peter RR White. “Appraisal theory”. In: *The international encyclopedia of language and social interaction* (2015), pp. 1–7.
- [95] Joachim Biermann et al. “Multi-level fusion of hard and soft information”. In: *17th International Conference on Information Fusion (FUSION)*. IEEE. 2014, pp. 1–8.
- [96] Erik Blasch. “Trust metrics in information fusion”. In: *Machine Intelligence and Bio-inspired Computation: Theory and Applications VIII*. Vol. 9119. International Society for Optics and Photonics. 2014, p. 91190L.
- [97] Erik Blasch et al. “URREF self-confidence in information fusion trust”. In: *Information Fusion (FUSION), 2014 17th International Conference on*. IEEE. 2014, pp. 1–8.
- [98] Valentina Dragos. *Evaluation of ONERA contributions to I2C : identification of threats and information retrieval*. RT 2/17417 - DTIM - TOULOUSE. 2014.

- [99] Souleiman Hasan and Edward Curry. “Approximate semantic matching of events for the internet of things”. In: *ACM Transactions on Internet Technology (TOIT)* 14.1 (2014), pp. 1–23.
- [100] Fabio Paglieri et al. “Trusting the messenger because of the message: feedback dynamics from information quality to source evaluation”. In: *Computational and Mathematical Organization Theory* 20.2 (2014), pp. 176–194.
- [101] Jeffrey Pennington, Richard Socher, and Christopher D Manning. “Glove: Global vectors for word representation”. In: *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 2014, pp. 1532–1543.
- [102] Fernando Roda and Estanislao Musulin. “An ontology-based framework to support intelligent data analysis of sensor measurements”. In: *Expert Systems with Applications* 41.17 (2014), pp. 7914–7926.
- [103] Erik Blasch et al. “URREF reliability versus credibility in information fusion (STANAG 2511)”. In: *Information Fusion (FUSION), 2013 16th International Conference on*. IEEE. 2013, pp. 1600–1607.
- [104] Anne-Claire Boury-Brisset. “Managing Semantic Big Data for Intelligence.” In: *STIDS*. 2013, pp. 41–47.
- [105] Marine Damiani and Delphine Battistelli. “Enunciative and modal variations in newswire texts in French: From guideline to automatic annotation”. In: *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*. 2013, pp. 223–227.
- [106] Valentina Dragos. *Construction d’une ontologie formelle pour l’aide au renseignement*. Séminaire TIM (Traitement de l’Information Multimodale) organisé par la DGA, Paris. June 2013.
- [107] Valentina Dragos. “Developing a core ontology to improve military intelligence analysis”. In: *International Journal of Knowledge-based and Intelligent Engineering Systems* 17.1 (2013), pp. 29–36.
- [108] Pek Hui Foo and Gee Wah Ng. “High-level information fusion: An overview.” In: *J. Adv. Inf. Fusion* 8.1 (2013), pp. 33–72.
- [109] Krzysztof Janowicz et al. “A restful proxy and data model for linked sensor data”. In: *International Journal of Digital Earth* 6.3 (2013), pp. 233–254.
- [110] Bahador Khaleghi et al. “Multisensor data fusion: A review of the state-of-the-art”. In: *Information Fusion* 14.1 (2013), pp. 28–44.
- [111] Philippe Lemercier. “The fundamentals of intelligence”. In: *Information Evaluation* (2013), pp. 55–102.
- [112] Christian Malewski et al. “Semantic matchmaking & mediation for sensors on the sensor web”. In: *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 7.3 (2013), pp. 929–934.
- [113] Kellyn Rein and Joachim Biermann. “Your high-level information is my low-level data-A new look at terminology for multi-level fusion”. In: *Proceedings of the 16th International Conference on Information Fusion*. IEEE. 2013, pp. 412–417.

- [114] Amit Sheth, Pramod Anantharam, and Cory Henson. “Physical-cyber-social computing: An early 21st century approach”. In: *IEEE Intelligent Systems* 1 (2013), pp. 78–82.
- [115] Erik Blasch et al. “Top ten trends in high-level information fusion”. In: *2012 15th International Conference on Information Fusion*. IEEE. 2012, pp. 2323–2330.
- [116] Paulo CG Costa et al. “Towards unbiased evaluation of uncertainty reasoning: The URREF ontology”. In: *2012 15th International Conference on Information Fusion*. IEEE. 2012, pp. 2301–2308.
- [117] Valentina Dragos. “Shallow semantic analysis to estimate HUMINT correlation”. In: *2012 15th International Conference on Information Fusion*. IEEE. 2012, pp. 2293–2300.
- [118] Jeffrey A Friedman and Richard Zeckhauser. “Assessing uncertainty in intelligence”. In: *Intelligence and National Security* 27.6 (2012), pp. 824–847.
- [119] Cory Henson, Amit Sheth, and Krishnaprasad Thirunarayan. “Semantic perception: Converting sensory observations to abstractions”. In: *IEEE Internet Computing* 16.2 (2012), pp. 26–34.
- [120] Christopher Soo-Guan Khoo, Armineh Nourbakhsh, and Jin-Cheon Na. “Sentiment analysis of online news text: a case study of appraisal theory”. In: *Online Information Review* (2012).
- [121] Mani Srivastava, Tarek Abdelzaher, and Boleslaw Szymanski. “Human-centric sensing”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 370.1958 (2012), pp. 176–197.
- [122] Robert H Dolin and Liora Alschuler. “Approaching semantic interoperability in health level seven”. In: *Journal of the American Medical Informatics Association* 18.1 (2011), pp. 99–103.
- [123] Hatem Ghorbel and David Jacot. “Sentiment analysis of French movie reviews”. In: *Advances in Distributed Agent-Based Retrieval Tools*. Springer, 2011, pp. 97–108.
- [124] Anne-Laure Jousselme et al. “Same world, different words: augmenting sensor output through semantics”. In: *14th International Conference on Information Fusion*. IEEE. 2011, pp. 127–135.
- [125] M El-Beze, A Jackiewicz, and S Hunston. “Opinions, sentiments et jugements d’évaluation”. In: *Traitement Automatique des Langues* 51.3 (2010), pp. 7–17.
- [126] Erik Blasch, Pierre Valin, and Eloi Bosse. “Measures of effectiveness for high-level fusion”. In: *2010 13th International Conference on Information Fusion*. IEEE. 2010, pp. 1–8.
- [127] Erik Blasch et al. “High level information fusion developments, issues, and grand challenges: Fusion 2010 panel discussion”. In: *2010 13th International Conference on Information Fusion*. IEEE. 2010, pp. 1–8.

- [128] Jordan Boyd-Graber and Philip Resnik. “Holistic sentiment analysis across languages: Multilingual supervised latent Dirichlet allocation”. In: *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. 2010, pp. 45–55.
- [129] David Lee Hall and John M Jordan. *Human-centered information fusion*. Artech House, 2010.
- [130] Galina L Rogova and Eloi Bosse. “Information quality in information fusion”. In: *2010 13th International Conference on Information Fusion*. IEEE. 2010, pp. 1–8.
- [131] Delphine Battistelli and Florence Amardeilh. “Knowledge Claims in Scientific Literature, Uncertainty and Semantic Annotation: A Case Study in the Biological Domain”. In: *Proceedings of the Semantic Authoring, Annotation and Knowledge Markup Workshop (SAAKM 2009)*. Los Angeles, United States, 2009.
- [132] Annie Piolat and Rachid Bannour. “EMOTAIX: un scénario de Tropes pour l’identification automatisée du lexique émotionnel et affectif”. In: *L’Année psychologique* 109.4 (2009), pp. 655–698.
- [133] Marco A Pravia et al. “Lessons learned in the creation of a data set for hard/soft information fusion”. In: *2009 12th International Conference on Information Fusion*. IEEE. 2009, pp. 2114–2121.
- [134] Kellyn Rein and Ulrich Schade. “How certain is certain? Evaluation of uncertainty in the fusion of information derived from diverse sources”. In: *2009 12th International Conference on Information Fusion*. IEEE. 2009, pp. 1613–1620.
- [135] Matthieu Vernier et al. “Catégorisation des évaluations dans un corpus de blogs multi-domaine”. In: *Revue des Nouvelles Technologies de l’Information* (2009), pp. 45–70.
- [136] Alain Auger and Jean Roy. “Expression of uncertainty in linguistic data”. In: *2008 11th International Conference on Information Fusion*. IEEE. 2008, pp. 1–8.
- [137] Rachel F Kesselman. “Verbal probability expressions in national intelligence estimates: a comprehensive analysis of trends from the fifties through post 9/11”. PhD thesis. MERCYHURST COLLEGE, 2008.
- [138] Nick R Milton. *Knowledge technologies*. Vol. 3. Polimetrica sas, 2008.
- [139] Soo Young Rieh and David R Danielson. “Credibility: A multidisciplinary framework”. In: *Annual review of information science and technology* 41.1 (2007), pp. 307–364.
- [140] Lauro Snidaro et al. “Quality-based fusion of multiple video sensors for video surveillance”. In: *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 37.4 (2007), pp. 1044–1051.
- [141] Óscar Ferrández et al. “Deep vs. Shallow semantic analysis applied to textual entailment recognition”. In: *International Conference on Natural Language Processing (in Finland)*. Springer. 2006, pp. 225–236.

- [142] Servane Gey and Jean-Michel Poggi. “Boosting and instability for regression trees”. In: *Computational statistics & data analysis* 50.2 (2006), pp. 533–550.
- [143] Asuncion Gomez-Perez, Mariano Fernández-López, and Oscar Corcho. *Ontological Engineering: with examples from the areas of Knowledge Management, e-Commerce and the Semantic Web*. Springer Science & Business Media, 2006.
- [144] Tom Gruber. “Every ontology is a treaty—a social agreement—among people with some common motive in sharing”. In: *Interview by Dr. Miltiadis D. Lytras, Official Quarterly Bulletin of AIS Special Interest Group on Semantic Web and Information Systems* 1.3 (2004).
- [145] Michael K Kalandros et al. “Tutorial on multisensor management and fusion algorithms for target tracking”. In: *Proceedings of the 2004 American control conference*. Vol. 5. IEEE. 2004, pp. 4734–4748.
- [146] Deborah L McGuinness, Frank Van Harmelen, et al. “OWL web ontology language overview”. In: *W3C recommendation* 10.10 (2004), p. 2004.
- [147] Xiao Hang Wang et al. “Ontology based context modeling and reasoning using OWL”. In: *IEEE annual conference on pervasive computing and communications workshops, 2004. Proceedings of the second*. Ieee. 2004, pp. 18–22.
- [148] Cristiano Castelfranchi, Rino Falcone, and Giovanni Pezzulo. “Trust in information sources as a source for trust: a fuzzy approach”. In: *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*. ACM. 2003, pp. 89–96.
- [149] Christopher J Matheus, Mieczyslaw M Kokar, and Kenneth Baclawski. “A core ontology for situation awareness”. In: *Proceedings of the Sixth International Conference on Information Fusion*. Vol. 1. 2003, pp. 545–552.
- [150] NATO. *STANAG 2011*. Tech. rep. Intelligence reports, 2003.
- [151] NATO. *The military intelligence data exchange standard - AIntP-3(B)*, NATO. Apr. 2003.
- [152] Edward Waltz. *Knowledge management in the intelligence enterprise*. Artech House, 2003.
- [153] Cristiano Castelfranchi and Rino Falcone. “Trust is much more than subjective probability: Mental components and sources of trust”. In: *System Sciences, 2000. Proceedings of the 33rd Annual Hawaii International Conference on*. IEEE. 2000, 10–pp.
- [154] John T Nockleby. “Hate speech”. In: *Encyclopedia of the American constitution* 3.2 (2000), pp. 1277–1279.
- [155] Graham JW King et al. “A standardized method for assessment of elbow function”. In: *Journal of shoulder and elbow surgery* 8.4 (1999), pp. 351–354.
- [156] Claudia Leacock and Martin Chodorow. “Combining local context and WordNet similarity for word sense identification”. In: *WordNet: An electronic lexical database* 49.2 (1998), pp. 265–283.

- [157] George A Miller. *WordNet: An electronic lexical database*. MIT press, 1998.
- [158] David L Hall and James Llinas. “An introduction to multisensor data fusion”. In: *Proceedings of the IEEE* 85.1 (1997), pp. 6–23.
- [159] Martin Ester et al. “A density-based algorithm for discovering clusters in large spatial databases with noise.” In: *Kdd*. Vol. 96. 34. 1996, pp. 226–231.
- [160] Thomas R Gruber. “A translation approach to portable ontology specifications”. In: *Knowledge acquisition* 5.2 (1993), pp. 199–220.
- [161] Lee Cronk. “Communication as manipulation: Implications for biosociological research”. In: *Proceedings of the American Sociological Association Annual Meetings*. 1991.
- [162] Sherman Kent. “Words of estimative probability”. In: *Studies in Intelligence* 8.4 (1964), pp. 49–65.
- [163] Charles Kay Ogden and Ivor Armstrong Richards. *The Meaning of Meaning: A Study of the Influence of Language upon Thought and of the Science of Symbolism*. Vol. 29. K. Paul, Trench, Trubner & Company, Limited, 1923.